

Alibaba Cloud Apsara Stack Enterprise Operations Guide

Version: 1909, Internal: V3.8.1

Issue: 20200116

Legal disclaimer

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

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequent

ial, exemplary, incidental, special, or punitive damages, including lost profits arising from the use or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

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

.

Document conventions

Style	Description	Example
	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: If the weight is set to 0, the server no longer receives new requests.
	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 <code>cd /d C:/window</code> command to enter the Windows system folder.
<i>Italic</i>	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid</code> <code>Instance_ID</code>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>

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

Contents

Legal disclaimer	I
Document conventions	I
1 Operations of basic platforms	1
1.1 Apsara Stack Operations (ASO)	1
1.1.1 Apsara Stack Operations overview	1
1.1.2 Log on to Apsara Stack Operations	3
1.1.3 Web page introduction	5
1.1.4 Operation and maintenance dashboard	6
1.1.5 Alarm Monitoring	6
1.1.5.1 Overview	6
1.1.5.2 Alert events	7
1.1.5.3 Alert history	12
1.1.5.4 Alert configuration	12
1.1.5.4.1 Alert contacts	12
1.1.5.4.2 Alert contact groups	13
1.1.5.4.3 Static parameter settings	14
1.1.5.5 Alert overview	15
1.1.5.6 Alert subscription and push	16
1.1.6 Resource Management	18
1.1.6.1 Physical servers	18
1.1.7 Inventory Management	20
1.1.7.1 View the ECS inventory	20
1.1.7.2 View the SLB inventory	29
1.1.7.3 View the RDS inventory	29
1.1.7.4 View the OSS inventory	30
1.1.7.5 View the Table Store inventory	30
1.1.7.6 View the Log Service inventory	31
1.1.7.7 View the EBS inventory	32
1.1.7.8 View the NAS inventory	32
1.1.7.9 View the HDFS inventory	33
1.1.8 Products	33
1.1.9 ITIL Management	34
1.1.9.1 Overview	34
1.1.9.2 Dashboard	35
1.1.9.3 Services	35
1.1.9.3.1 Basic functions	35
1.1.9.3.1.1 Overview	35
1.1.9.3.1.2 Manage requests	36
1.1.9.3.1.3 Manage tasks	37
1.1.9.3.2 Manage incidents	38

1.1.9.3.2.1 Create an incident request.....	38
1.1.9.3.2.2 Manage incident requests.....	40
1.1.9.3.2.3 Manage incident tasks.....	42
1.1.9.3.3 Manage problems.....	44
1.1.9.3.3.1 Create a problem request.....	44
1.1.9.3.3.2 Manage problem requests.....	45
1.1.9.3.3.3 Manage problem tasks.....	47
1.1.9.4 Version control.....	49
1.1.9.5 Configure process templates.....	50
1.1.9.6 Configure CAB or ECAB.....	53
1.1.10 Configurations.....	54
1.1.10.1 Overview.....	54
1.1.10.2 Modify a configuration item of a product.....	54
1.1.10.3 Restore the configuration value of a modified configuration item.....	55
1.1.10.4 Manage kernel configurations.....	56
1.1.10.5 Scan configurations.....	56
1.1.11 Offline Backup.....	57
1.1.11.1 Service configuration.....	57
1.1.11.1.1 Configure the backup server.....	57
1.1.11.1.2 Add a backup product.....	58
1.1.11.2 Backup service.....	59
1.1.11.2.1 Backup configuration.....	59
1.1.11.2.2 View the backup details.....	61
1.1.11.2.3 View the backup server status.....	61
1.1.11.3 View the backup status.....	61
1.1.12 NOC.....	62
1.1.12.1 Overview.....	62
1.1.12.2 Dashboard.....	62
1.1.12.3 Network topology.....	64
1.1.12.4 Resource management.....	64
1.1.12.4.1 Network elements.....	64
1.1.12.4.1.1 Device management.....	65
1.1.12.4.1.2 Modify the device password.....	69
1.1.12.4.1.3 Configuration comparison.....	70
1.1.12.4.2 Server Load Balancers.....	71
1.1.12.4.2.1 View the cluster monitoring information.....	71
1.1.12.4.2.2 View the instance monitoring information.....	71
1.1.12.4.3 Collect IP addresses.....	72
1.1.12.4.4 IP address ranges.....	73
1.1.12.4.4.1 Import the planning file.....	73
1.1.12.4.4.2 Manually add the IP address pool information.....	73
1.1.12.4.4.3 Modify the IP address pool information.....	74
1.1.12.4.4.4 Export the IP address pool information.....	74
1.1.12.4.4.5 Delete the IP address pool information.....	74

1.1.12.5 Alert management.....	75
1.1.12.5.1 View and process current alerts.....	75
1.1.12.5.2 View history alerts.....	76
1.1.12.5.3 Add a trap.....	76
1.1.12.5.4 View a trap.....	79
1.1.12.6 Network reconfiguration.....	79
1.1.12.6.1 Physical network integration.....	79
1.1.12.6.2 ASW scale-up.....	81
1.1.12.7 Fault check.....	84
1.1.12.7.1 IP address conflict check.....	84
1.1.12.7.2 Leased line discovery.....	84
1.1.12.7.3 Network inspection.....	86
1.1.12.7.4 Configuration baseline audit.....	88
1.1.13 Full Stack Monitor.....	89
1.1.13.1 SLA.....	89
1.1.13.1.1 View the current state of a cloud product.....	89
1.1.13.1.2 View the history data of a cloud product.....	89
1.1.13.1.3 View the availability of an instance.....	90
1.1.13.1.4 View the availability of a product.....	90
1.1.13.2 Operations full link logs.....	91
1.1.13.3 Correlation diagnosis and alarm.....	92
1.1.13.3.1 Full stack correlation alert.....	92
1.1.13.3.2 Server.....	93
1.1.13.3.3 Network device.....	96
1.1.13.3.4 ECS.....	96
1.1.13.3.5 RDS.....	98
1.1.13.3.6 SLB.....	99
1.1.13.3.7 VPC.....	100
1.1.14 Pangu Operation Center.....	101
1.1.14.1 Pangu grail.....	101
1.1.14.2 Cluster information.....	102
1.1.14.3 Node information.....	103
1.1.14.4 Pangu operation.....	104
1.1.15 Task Management.....	105
1.1.15.1 Overview.....	105
1.1.15.2 View the task overview.....	105
1.1.15.3 Create a script.....	106
1.1.15.4 Create a task.....	107
1.1.15.5 Modify a script.....	108
1.1.15.6 Delete a script.....	109
1.1.15.7 View the execution status of a task.....	109
1.1.15.8 Modify the execution parameters of a task.....	109
1.1.15.9 Start a task.....	110
1.1.15.10 Delete a task.....	110
1.1.16 Process tasks to be intervened.....	111

1.1.17 System Management.....	111
1.1.17.1 Overview.....	111
1.1.17.2 Department management.....	111
1.1.17.3 Role management.....	112
1.1.17.4 Logon policy management.....	113
1.1.17.5 User management.....	114
1.1.17.6 Two factor authentication.....	117
1.1.17.7 Application whitelist.....	121
1.1.17.8 Server password management.....	122
1.1.17.9 Operation logs.....	124
1.1.17.10 View the authorization information.....	125
1.2 Apsara Stack Doctor (ASD).....	127
1.2.1 Apsara Stack Doctor introduction.....	127
1.2.2 Log on to Apsara Stack Doctor.....	130
1.2.3 Product dependency.....	132
1.2.4 Apsara Stack Inspection System.....	135
1.2.4.1 Apsara Stack Inspection System introduction.....	135
1.2.4.2 Log on to Apsara Stack Inspection System.....	136
1.2.4.3 Apsara Stack Inspection System overview.....	137
1.2.4.4 Platform Inspection.....	140
1.2.4.4.1 Overview.....	140
1.2.4.4.2 Basic Inspection.....	140
1.2.4.4.3 Apsara System Inspection.....	141
1.2.4.4.4 Inspection in Other Systems.....	142
1.2.4.4.5 Inventory Inspection.....	143
1.2.4.4.6 Cloud Product Inspection.....	145
1.2.4.4.7 Middleware Inspection.....	145
1.2.4.4.8 Big Data Inspection.....	146
1.2.4.4.9 Inspection Reports.....	146
1.2.4.5 Inspection History.....	147
1.2.4.6 Work Reports.....	148
1.2.4.6.1 Add project information.....	148
1.2.4.6.2 Issues.....	149
1.2.4.6.2.1 Add issues.....	149
1.2.4.6.2.2 Update the issue handling progress.....	149
1.2.4.6.2.3 Remove issues.....	150
1.2.4.6.3 View resource utilization information.....	150
1.2.4.6.4 Reports.....	151
1.2.4.6.4.1 Add issues.....	151
1.2.4.6.4.2 Update issue progress.....	151
1.2.4.6.4.3 Remove issues.....	152
1.2.4.6.4.4 View unresolved issues.....	152
1.2.4.6.4.5 View reports.....	152
1.2.4.6.5 Generate summary reports.....	152
1.2.4.6.6 Download work reports.....	153

1.2.4.7 Products.....	154
1.2.4.8 E2E ECS Links.....	154
1.2.5 ASA.....	155
1.2.5.1 RPM Check.....	155
1.2.5.2 Virtual IP Check.....	156
1.2.5.3 Volume Check.....	157
1.2.5.4 NTP Check.....	158
1.2.5.5 IP Conflict Check.....	159
1.2.5.6 DNS Check.....	160
1.2.5.7 IP Details.....	161
1.2.5.8 Quota Check.....	161
1.2.5.9 Error Diagnostics.....	162
1.2.5.10 Versions.....	163
1.2.6 Support tools.....	163
1.2.6.1 Diagnose with the OS tool.....	163
1.2.6.2 Use Support Tools.....	164
1.2.6.3 Update Support Tools.....	166
1.2.6.4 Diagnose with inspection tools.....	167
1.2.6.5 Upload script files for EDAS diagnostics.....	168
1.2.6.6 EDAS diagnostics.....	169
1.2.7 Service Availability.....	170
1.2.7.1 View Service Availability.....	170
1.3 Operation Access Manager (OAM).....	171
1.3.1 OAM introduction.....	171
1.3.2 Instructions.....	172
1.3.3 Quick start.....	174
1.3.3.1 Log on to OAM.....	174
1.3.3.2 Create a group.....	175
1.3.3.3 Add group members.....	175
1.3.3.4 Add group roles.....	176
1.3.3.5 Create a role.....	177
1.3.3.6 Add inherited roles to a role.....	177
1.3.3.7 Add resources to a role.....	177
1.3.3.8 Add authorized users to a role.....	178
1.3.4 Manage groups.....	180
1.3.4.1 Modify the group information.....	180
1.3.4.2 View group role details.....	180
1.3.4.3 Delete a group.....	181
1.3.4.4 View authorized groups.....	181
1.3.5 Manage roles.....	182
1.3.5.1 Search for roles.....	182
1.3.5.2 Modify the role information.....	182
1.3.5.3 View the role inheritance tree.....	182
1.3.5.4 Transfer roles.....	183
1.3.5.5 Delete a role.....	183

1.3.5.6 View authorized roles.....	184
1.3.5.7 View all roles.....	184
1.3.6 Search for resources.....	184
1.3.7 View the personal information.....	185
1.3.8 Appendix.....	185
1.3.8.1 Default roles and their functions.....	185
1.3.8.1.1 Default role of OAM.....	185
1.3.8.1.2 Default roles of Apsara Infrastructure Management Framework.....	186
1.3.8.1.3 Default role of DataQ - Smart Tag Service.....	188
1.3.8.1.4 Default roles of Webapp-rule.....	189
1.3.8.1.5 Default roles of the workflow console.....	190
1.3.8.1.6 Default role of Tianjimon.....	190
1.3.8.2 Permission lists of operations platforms.....	191
1.3.8.2.1 Permission list of Apsara Infrastructure Management Framework.....	191
1.3.8.2.2 Permission list of DataQ - Smart Tag Service.....	201
1.3.8.2.3 Permission list of Webapp-rule.....	201
1.3.8.2.4 Permission list of the workflow console.....	201
1.3.8.2.5 Permission list of Tianjimon.....	202
1.4 Apsara Opsapi Management.....	202
1.4.1 Apsara Opsapi Management system overview.....	202
1.4.2 Log on to the Apsara Opsapi Management system.....	203
1.4.3 API management.....	205
1.4.3.1 Register APIs.....	205
1.4.3.2 Modify information about APIs.....	205
1.4.3.3 Test APIs.....	207
1.4.3.4 Remove information about APIs.....	207
1.4.3.5 API design.....	208
1.4.3.5.1 Designers.....	208
1.4.3.5.2 Designer nodes.....	208
1.4.3.5.3 Design an API flow.....	209
1.4.4 Version management.....	210
1.4.4.1 Apsara Stack version management.....	210
1.4.4.1.1 Add information about versions.....	210
1.4.4.1.2 Select products for an Apsara Stack version.....	210
1.4.4.1.3 Compare versions.....	211
1.4.4.1.4 Remove information about Apsara Stack versions.....	214
1.4.4.2 Product baseline management.....	214
1.4.4.3 Product management.....	215
1.4.4.3.1 Add information about products.....	215
1.4.4.3.2 Add information about product versions.....	216
1.4.4.3.3 Import information about APIs.....	216
1.4.4.3.4 Set SDK versions.....	217
1.4.4.3.5 Modify product names and descriptions.....	218

1.4.4.3.6 View information about product versions.....	218
1.4.4.3.7 Modify information about product versions.....	218
1.4.4.3.8 Remove information about product versions.....	219
1.4.4.3.9 Remove information about products.....	219
1.4.4.3.10 Remove information about product APIs.....	219
1.4.4.4 SDK management.....	220
1.4.4.4.1 Customize SDKs.....	220
1.4.4.4.2 Modify SDKs.....	221
1.4.4.4.3 Delete SDKs.....	222
1.4.5 Test management.....	222
1.4.5.1 Test cases.....	223
1.4.5.1.1 Modify test cases.....	223
1.4.5.1.2 Run test cases.....	224
1.4.5.1.3 Delete test cases.....	225
1.4.5.2 Test sets.....	225
1.4.5.2.1 Create test sets.....	225
1.4.5.2.2 Associate test cases.....	225
1.4.5.2.3 Run test sets.....	226
1.4.5.2.4 Delete test sets.....	226
1.4.6 View execution history of test cases.....	227
1.4.7 System management.....	227
1.4.7.1 Metadatabase management.....	227
1.4.7.1.1 View information about added metadatabases.....	227
1.4.7.1.2 View connection information about metadatabases.....	229
1.4.7.1.3 Remove information about metadatabases.....	229
1.4.7.2 Server management.....	229
1.4.7.2.1 View information about added servers.....	229
1.4.7.2.2 Remove server information.....	231
1.4.7.3 Audit APIs.....	232
1.4.7.4 View logs.....	232
1.5 Apsara Infrastructure Management Framework.....	233
1.5.1 What is Apsara Infrastructure Management Framework?.....	233
1.5.1.1 Overview.....	233
1.5.1.2 Basic concepts.....	234
1.5.2 Log on to Apsara Infrastructure Management Framework.....	236
1.5.3 Web page introduction.....	238
1.5.3.1 Introduction on the home page.....	238
1.5.3.2 Introduction on the left-side navigation pane.....	241
1.5.4 Cluster operations.....	243
1.5.4.1 View cluster configurations.....	243
1.5.4.2 View the cluster dashboard.....	245
1.5.4.3 View the cluster operation and maintenance center.....	250
1.5.4.4 View the service final status.....	254
1.5.4.5 View operation logs.....	256
1.5.5 Service operations.....	257

1.5.5.1 View the service list.....	257
1.5.5.2 View the service instance dashboard.....	258
1.5.5.3 View the server role dashboard.....	261
1.5.6 Machine operations.....	264
1.5.6.1 View the machine dashboard.....	264
1.5.7 Monitoring center.....	267
1.5.7.1 Modify an alert rule.....	267
1.5.7.2 View the status of a monitoring instance.....	267
1.5.7.3 View the alert status.....	268
1.5.7.4 View alert rules.....	268
1.5.7.5 View the alert history.....	269
1.5.8 Tasks and deployment summary.....	270
1.5.8.1 View rolling tasks.....	270
1.5.8.2 View running tasks.....	272
1.5.8.3 View history tasks.....	273
1.5.8.4 View the deployment summary.....	273
1.5.9 Reports.....	276
1.5.9.1 View reports.....	276
1.5.9.2 Add a report to favorites.....	278
1.5.10 Appendix.....	278
1.5.10.1 Project component info report.....	278
1.5.10.2 IP list.....	279
1.5.10.3 Machine info report.....	279
1.5.10.4 Rolling info report.....	281
1.5.10.5 Machine RMA approval pending list.....	283
1.5.10.6 Registration vars of services.....	285
1.5.10.7 Virtual machine mappings.....	285
1.5.10.8 Service inspector report.....	285
1.5.10.9 Resource application report.....	286
1.5.10.10 Statuses of project components.....	287
1.5.10.11 Relationship of service dependency.....	289
1.5.10.12 Check report of network topology.....	289
1.5.10.13 Clone report of machines.....	290
1.5.10.14 Auto healing/install approval pending report.....	291
1.5.10.15 Machine power on or off statuses of clusters.....	291
1.6 Network operations.....	292
1.6.1 Apsara Network Intelligence.....	292
1.6.1.1 What is Apsara Network Intelligence?.....	293
1.6.1.2 Log on to the Apsara Network Intelligence console.....	293
1.6.1.3 Query information.....	293
1.6.1.4 Manage cloud service instances.....	295
1.6.1.5 Tunnel VIP.....	295
1.6.1.5.1 Create a Layer-4 listener VIP.....	295
1.6.1.5.2 Query the tunnel VIP of a cloud service.....	296
1.6.1.6 Create a Direct Any Tunnel VIP.....	297

1.6.1.7 Leased line connection.....	297
1.6.1.7.1 Overview.....	297
1.6.1.7.2 Manage an access point.....	298
1.6.1.7.3 Manage an access device.....	299
1.6.1.7.4 Establish a leased line connection.....	301
1.6.1.7.5 Create a VBR.....	305
1.6.1.7.6 Create router interfaces.....	308
1.6.1.7.7 Create a routing table.....	311
1.6.1.8 Manage Business Foundation System flows in a VPC.....	313
1.6.1.9 Configure reverse access to cloud services.....	314
2 Operations of basic cloud products.....	316
2.1 Elastic Compute Service (ECS).....	316
2.1.1 ECS overview.....	316
2.1.2 Log on to Apsara Stack Operations.....	317
2.1.3 ECS operations and maintenance.....	319
2.1.3.1 Overview.....	319
2.1.3.2 VM.....	319
2.1.3.2.1 Overview.....	319
2.1.3.2.2 Search for VMs.....	319
2.1.3.2.3 Start a VM.....	319
2.1.3.2.4 Stop a VM.....	320
2.1.3.2.5 Restart a VM.....	321
2.1.3.2.6 Cold migration.....	321
2.1.3.2.7 Reset a disk.....	322
2.1.3.3 Disks.....	323
2.1.3.3.1 Overview.....	323
2.1.3.3.2 Search for disks.....	323
2.1.3.3.3 Search for snapshots.....	323
2.1.3.3.4 Mount a disk.....	324
2.1.3.3.5 Detach a disk.....	324
2.1.3.3.6 Create a snapshot.....	324
2.1.3.4 Snapshots.....	325
2.1.3.4.1 Overview.....	325
2.1.3.4.2 Search for snapshots.....	325
2.1.3.4.3 Delete a snapshot.....	326
2.1.3.4.4 Create an image.....	326
2.1.3.5 Images.....	326
2.1.3.5.1 Overview.....	326
2.1.3.5.2 Search for images.....	327
2.1.3.6 Security groups.....	327
2.1.3.6.1 Overview.....	327
2.1.3.6.2 Search for security groups.....	327
2.1.3.6.3 Add security group rules.....	328
2.1.4 VM hot migration.....	329
2.1.4.1 Overview.....	329

2.1.4.2 Limits on hot migration.....	330
2.1.4.3 Complete hot migration on AG.....	331
2.1.4.4 Modify the position of the NC where the VM is located.....	333
2.1.4.5 FAQ.....	333
2.1.5 Hot migration of disks.....	336
2.1.5.1 Overview.....	336
2.1.5.2 Limits.....	336
2.1.5.3 O&M after hot migration.....	337
2.1.6 Upgrade solution.....	337
2.1.6.1 Overview.....	337
2.1.6.2 Limits on GPU clusters.....	337
2.1.6.3 Limits on FPGA clusters.....	338
2.1.7 Disk maintenance of an instance.....	338
2.1.7.1 Overview.....	338
2.1.7.2 Maintenance procedure.....	339
2.1.7.3 Additional instructions.....	350
2.1.8 Handle routine alarms.....	351
2.1.8.1 Overview.....	351
2.1.8.2 API proxy.....	353
2.1.8.3 API Server.....	353
2.1.8.4 RegionMaster.....	354
2.1.8.5 RMS.....	355
2.1.8.6 PYNC.....	356
2.1.8.7 Zookeeper.....	357
2.1.8.8 AG.....	357
2.1.8.9 Server groups.....	358
2.1.9 Inspection.....	359
2.1.9.1 Overview.....	359
2.1.9.2 Cluster basic health inspection.....	359
2.1.9.2.1 Overview.....	359
2.1.9.2.2 Monitoring inspection.....	359
2.1.9.2.3 Inspection of basic software package versions.....	359
2.1.9.2.4 Basic public resources inspection.....	359
2.1.9.3 Cluster resource inspection.....	360
2.1.9.3.1 Overview.....	360
2.1.9.3.2 Cluster inventory inspection.....	360
2.1.9.3.3 VM inspection.....	362
2.2 Container Service.....	363
2.2.1 Components and features.....	363
2.2.1.1 Console.....	363
2.2.1.2 Troopers.....	364
2.2.1.3 Mirana.....	365
2.2.2 System restart.....	366
2.2.2.1 Restart a control node.....	366
2.2.3 Security maintenance.....	367

2.2.3.1 Network security maintenance.....	367
2.3 Auto Scaling (ESS).....	367
2.3.1 Log on to Apsara Stack Operations.....	367
2.3.2 Product resources and services.....	369
2.3.2.1 Application deployment.....	369
2.3.2.2 Troubleshooting.....	369
2.3.3 Inspection.....	370
2.3.3.1 Overview.....	370
2.3.3.2 Monitoring inspection.....	371
2.3.3.3 Basic software package version inspection.....	371
2.4 Object Storage Service (OSS).....	371
2.4.1 Log on to the Apsara Stack Operations console.....	371
2.4.2 OSS operations and maintenance.....	372
2.4.2.1 User data.....	373
2.4.2.1.1 Basic bucket information.....	373
2.4.2.1.2 User data overview.....	373
2.4.2.1.3 Data monitoring.....	374
2.4.2.2 Cluster data.....	376
2.4.2.2.1 Inventory monitoring.....	376
2.4.2.2.2 Bucket statistics.....	377
2.4.2.2.3 Object statistics.....	378
2.4.2.2.4 Data monitoring.....	379
2.4.2.2.5 Resource usage rankings.....	382
2.4.3 Use of tools.....	383
2.4.3.1 Typical commands supported by tsar.....	383
2.4.3.2 Configure tsar for statistic collection.....	383
2.5 Table Store.....	383
2.5.1 Table Store Operations and Maintenance System.....	383
2.5.1.1 Overview.....	383
2.5.1.2 User data.....	384
2.5.1.2.1 Instance management.....	384
2.5.1.3 Cluster management.....	388
2.5.1.3.1 Cluster information.....	388
2.5.1.4 Inspection center.....	390
2.5.1.4.1 Abnormal resource usage.....	390
2.5.1.5 Monitoring center.....	391
2.5.1.5.1 Cluster monitoring.....	391
2.5.1.5.2 Application monitoring.....	392
2.5.1.5.3 Top requests.....	393
2.5.1.5.4 Request log search.....	394
2.5.1.6 System management.....	394
2.5.1.6.1 Manage tasks.....	394
2.5.1.6.2 View tasks.....	396
2.5.1.7 Platform audit.....	396
2.5.1.7.1 Operation logs.....	396

2.5.2 Cluster environments.....	397
2.5.3 System roles.....	398
2.5.4 Pre-partition a table.....	399
2.5.4.1 Pre-partitioning.....	399
2.5.4.2 View partitions.....	400
2.6 Apsara File Storage for HDFS.....	401
2.6.1 Overview.....	401
2.6.1.1 Overview of Apsara Distributed File System.....	401
2.6.1.2 Overview of Apsara Infrastructure Management Framework...	403
2.6.1.2.1 Terms of Apsara Infrastructure Management Framework....	403
2.6.2 Configuration update.....	404
2.6.2.1 Overview.....	404
2.6.2.2 Procedure.....	405
2.6.2.3 Configure the file structure.....	407
2.6.2.4 Make the configuration changes take effect.....	409
2.6.2.5 Overwrite configurations.....	409
2.6.2.6 Configure validity checks.....	410
2.6.3 Cluster operations and maintenance in Apsara Distributed File System.....	410
2.6.3.1 Set a global flag in Apsara Distributed File System.....	410
2.6.3.2 Manage files in Apsara Distributed File System.....	411
2.6.3.3 Common puadmin commands.....	412
2.6.3.4 GC of Apsara Distributed File System.....	413
2.6.3.5 Cluster rebalance.....	414
2.6.3.6 Directory quota operations.....	416
2.6.3.7 Directory pin operations.....	418
2.6.4 Operations and maintenance of masters.....	419
2.6.4.1 Overview.....	419
2.6.4.2 Switch over the primary master.....	420
2.6.4.3 Status check for multiple masters.....	421
2.6.4.3.1 View the election status.....	421
2.6.4.3.2 View the log synchronization status of multiple masters.....	421
2.6.4.4 Rules for the Apsara Distributed File System master to write .cpt and .log files.....	422
2.6.4.5 Master replacement.....	422
2.6.4.5.1 Procedure.....	422
2.6.4.5.2 Replace a master.....	425
2.6.4.5.3 Manually replace a master.....	426
2.6.4.6 Manually synchronize logs between a primary master and a secondary master.....	427
2.6.4.7 Rename a chunkserver online.....	428
2.6.4.8 Multi-master tools.....	429
2.6.5 Operations and maintenance of chunkservers.....	429
2.6.5.1 Set the chunkserver status.....	429
2.6.5.2 Set the disk status.....	430

2.6.5.3 Scale-out and scale-in of chunkservers.....	431
2.6.5.3.1 Procedure.....	431
2.6.5.3.2 Scale-out procedure.....	432
2.6.5.3.3 Scale-in procedure.....	432
2.6.5.3.4 Disable manual scale-out of chunkservers by using puadmin.....	434
2.6.6 Track the cluster status.....	435
2.6.6.1 View the cluster status.....	435
2.6.6.2 Submission history.....	435
2.6.7 FAQ.....	436
2.6.7.1 Identify the machine running pangu_supervisor and the log location.....	436
2.6.7.2 Adjust the flag of pangu_supervisor.....	437
2.6.7.3 Supervisor approval errors.....	437
2.6.7.3.1 Supervisor approval prerequisites.....	437
2.6.7.3.2 Failure to approve the master replacement.....	439
2.6.7.3.3 Failure to approve chunkserver disconnection.....	439
2.6.7.4 Accelerate chunkserver disconnection.....	440
2.6.7.5 Rolling failure during a hot upgrade of Apsara Distributed File System.....	440
2.6.7.6 Manually replace binary files in an emergency.....	440
2.6.7.6.1 Manual overwrite operation.....	440
2.6.7.6.2 Use the overwrite tool of Apsara Infrastructure Management Framework.....	442
2.7 ApsaraDB for RDS.....	442
2.7.1 Architecture.....	442
2.7.1.1 System architecture.....	442
2.7.1.1.1 Backup system.....	443
2.7.1.1.2 Data migration system.....	443
2.7.1.1.3 Monitoring system.....	444
2.7.1.1.4 Control system.....	445
2.7.1.1.5 Task scheduling system.....	446
2.7.2 RDS O&M overview.....	446
2.7.3 Log on to Apsara Stack Operations console.....	446
2.7.4 Manage instances.....	448
2.7.5 Manage hosts.....	450
2.7.6 Security maintenance.....	451
2.7.6.1 Network security maintenance.....	451
2.7.6.2 Account password maintenance.....	452
2.8 KVStore for Redis.....	452
2.8.1 O&M tool.....	452
2.8.2 Architecture diagram.....	452
2.8.3 Log on to Apsara Stack Operations.....	452
2.8.4 Instance management.....	454
2.8.5 Host management.....	454

2.8.6 Security maintenance.....	455
2.8.6.1 Network security maintenance.....	455
2.8.6.2 Password maintenance.....	456
2.9 ApsaraDB for MongoDB.....	456
2.9.1 Service architecture.....	456
2.9.1.1 System architecture.....	456
2.9.1.1.1 Backup system.....	456
2.9.1.1.2 Data migration system.....	457
2.9.1.1.3 Monitoring system.....	457
2.9.1.1.4 Control system.....	458
2.9.1.1.5 Task scheduling system.....	458
2.9.2 ApsaraDB for MongoDB O&M overview.....	458
2.9.3 Log on to the Apsara Stack Operations console.....	459
2.9.4 Manage instances.....	460
2.9.5 Host management.....	462
2.9.6 Security maintenance.....	463
2.9.6.1 Network security maintenance.....	463
2.9.6.2 Account password maintenance.....	464
2.10 AnalyticDB for PostgreSQL.....	464
2.10.1 Overview.....	464
2.10.2 System architecture.....	466
2.10.3 Routine maintenance.....	466
2.10.3.1 Check for data skew on a regular basis.....	467
2.10.3.2 Execute the VACUUM and ANALYZE statements.....	468
2.10.4 Security maintenance.....	468
2.10.4.1 Network security maintenance.....	468
2.10.4.2 Account password maintenance.....	469
2.11 Apsara Stack Security.....	469
2.11.1 Log on to the Apsara Infrastructure Management Framework console.....	469
2.11.2 Routine operations and maintenance of Server Guard.....	469
2.11.2.1 Check the service status.....	469
2.11.2.1.1 Check the client status.....	469
2.11.2.1.2 Check the status of Aegiserver.....	470
2.11.2.1.3 Check the Server Guard Update Service status.....	472
2.11.2.1.4 Check the Defender module status.....	472
2.11.2.2 Restart Server Guard.....	473
2.11.3 Routine operations and maintenance of Network Traffic Monitoring System.....	475
2.11.3.1 Check the service status.....	475
2.11.3.1.1 Basic inspection.....	475
2.11.3.1.2 Advanced inspection.....	475
2.11.3.2 Common operations and maintenance.....	477
2.11.3.2.1 Restart the Network Traffic Monitoring System process.....	477
2.11.3.2.2 Uninstall Network Traffic Monitoring System.....	477

2.11.3.2.3 Disable TCP blocking.....	478
2.11.3.2.4 Enable TCPDump.....	478
2.11.4 Routine operations and maintenance of Anti-DDoS Service.....	479
2.11.4.1 Check the service status.....	479
2.11.4.1.1 Basic inspection.....	479
2.11.4.1.2 Advanced inspection.....	479
2.11.4.2 Common operations and maintenance.....	482
2.11.4.2.1 Restart Anti-DDoS Service.....	482
2.11.4.2.2 Troubleshoot common faults.....	483
2.11.5 Routine operations and maintenance of Threat Detection Service.....	488
2.11.5.1 Check the service status.....	488
2.11.5.1.1 Basic inspection.....	488
2.11.5.1.2 Advanced inspection.....	488
2.11.5.2 Restart TDS.....	490
2.11.6 Routine operations and maintenance of WAF.....	490
2.11.6.1 Check the service status.....	490
2.11.6.1.1 Basic inspection.....	490
2.11.6.1.2 Advanced inspection.....	491
2.11.7 Routine operations and maintenance of Sensitive Data Discovery and Protection.....	494
2.11.7.1 Check the service status.....	494
2.11.7.1.1 Basic inspection.....	494
2.11.7.1.2 Advanced inspection: Check the status of the SddpService service.....	495
2.11.7.1.3 Advanced inspection: Check the status of the SddpData service.....	497
2.11.7.1.4 Advanced inspection: Check the status of the SddpPrivilege service.....	498
2.11.7.1.5 Advanced inspection: Check the status of the SddpLog service.....	500
2.11.7.2 Restart SDDP.....	501
2.11.8 Routine operations and maintenance of Apsara Stack Security Center.....	503
2.11.8.1 Check service status.....	503
2.11.8.1.1 Basic inspection.....	503
2.11.8.1.2 Advanced inspection.....	503
2.11.8.2 Restart the secure-console service.....	504
2.11.9 Routine operations and maintenance of secure-service.....	505
2.11.9.1 Check the service status.....	505
2.11.9.1.1 Basic inspection.....	505
2.11.9.1.2 Advanced inspection: Check the secure-service status.....	505
2.11.9.1.3 Check the Dolphin service status.....	507
2.11.9.1.4 Check the data-sync service status.....	508
2.11.9.2 Restart secure-service.....	508

2.12 Key Management Service (KMS)	510
2.12.1 Operations and maintenance of KMS components	510
2.12.1.1 Overview.....	510
2.12.1.2 KMS_HOST.....	510
2.12.1.3 HSA.....	515
2.12.1.4 etcd.....	520
2.12.1.5 Rotator.....	522
2.12.1.5.1 Primary IDC.....	522
2.12.1.5.2 Secondary IDC.....	523
2.12.2 Log analysis.....	524
2.12.2.1 Overview.....	524
2.12.2.2 Request IDs.....	525
2.12.2.3 Common KMS errors.....	526
2.12.2.3.1 Overview.....	526
2.12.2.3.2 Error codes 4xx.....	526
2.12.2.3.3 Error code 500.....	526
2.12.2.3.4 Error code 503.....	526
2.12.2.3.5 Dependent service degradation.....	527
2.12.3 View and process internal data.....	528
2.13 Log Service	530
2.13.1 Component O&M.....	530
2.13.2 Troubleshooting.....	532
2.13.2.1 NGINX.....	532
2.13.2.2 Console.....	532
2.13.2.3 Service.....	532
2.14 Apsara Stack DNS	532
2.14.1 Introduction to Apsara Stack DNS.....	533
2.14.2 Maintenance.....	533
2.14.2.1 View operational logs.....	533
2.14.2.2 Enable and disable a service.....	533
2.14.2.3 Data backup.....	534
2.14.3 DNS API.....	534
2.14.3.1 Manage the API system.....	534
2.14.3.2 Troubleshooting.....	537
2.14.4 DNS system.....	538
2.14.4.1 Check whether the service role is normal.....	538
2.14.4.2 Troubleshooting.....	540
2.14.4.3 Errors and exceptions.....	540
2.14.5 Log analysis.....	541
2.14.6 View and process data.....	541
2.15 API Gateway	541
2.15.1 API Gateway introduction.....	541
2.15.2 Routine maintenance.....	542
2.15.2.1 View operation logs.....	542
2.15.2.2 Enable and disable API services.....	542

2.15.3 API Gateway O&M.....	542
2.15.3.1 System O&M.....	542
2.15.3.1.1 Check the desired state of API Gateway.....	542
2.15.3.1.2 Check the service status of OpenAPI.....	543
2.15.3.1.3 Check the service status of the API Gateway console.....	545
2.15.3.1.4 Check the service status of API Gateway.....	547
2.15.3.1.5 View results of automated test cases.....	549
2.15.3.2 Troubleshooting.....	549
2.15.4 Log analysis.....	550
3 Operations of big data products.....	551
3.1 Apsara Bigdata Manager (ABM) platform.....	551
3.1.1 What is Apsara Bigdata Manager?.....	551
3.1.2 Common operations.....	551
3.1.3 Quick start.....	558
3.1.3.1 Log on to the ABM console.....	558
3.1.3.2 Set the background color.....	560
3.1.3.3 View the dashboard.....	562
3.1.3.4 View the cluster running status.....	567
3.1.3.5 View and clear cluster alerts.....	569
3.1.4 ABM.....	573
3.1.4.1 ABM dashboard.....	573
3.1.4.2 ABM repository.....	582
3.1.4.3 O&M overview.....	583
3.1.4.4 Cluster O&M.....	585
3.1.4.4.1 Cluster overview.....	585
3.1.4.4.2 Cluster health.....	589
3.1.4.5 Service O&M.....	595
3.1.4.5.1 Service overview.....	595
3.1.4.5.2 Service hosts.....	600
3.1.4.6 Host O&M.....	600
3.1.4.6.1 Host overview.....	600
3.1.4.6.2 Host health.....	606
3.1.5 MaxCompute.....	610
3.1.5.1 MaxCompute workbench.....	610
3.1.5.2 Business O&M.....	615
3.1.5.2.1 Business O&M overview.....	615
3.1.5.2.2 Project management.....	616
3.1.5.2.2.1 Project list.....	617
3.1.5.2.2.2 Disaster recovery.....	620
3.1.5.2.3 Job management.....	625
3.1.5.2.3.1 Job snapshots.....	625
3.1.5.2.4 Business optimization.....	628
3.1.5.2.4.1 File merging.....	628
3.1.5.2.4.2 File archiving.....	631
3.1.5.2.4.3 Resource analysis.....	634

3.1.5.3 Service O&M.....	636
3.1.5.3.1 Control service O&M.....	636
3.1.5.3.1.1 Control service O&M overview.....	636
3.1.5.3.1.2 Control service overview.....	637
3.1.5.3.1.3 Control service health.....	639
3.1.5.3.1.4 Control service instances.....	639
3.1.5.3.1.5 Disable or enable a control service role.....	640
3.1.5.3.1.6 Start AdminConsole.....	641
3.1.5.3.1.7 Collect service logs.....	642
3.1.5.3.2 Job Scheduler O&M.....	644
3.1.5.3.2.1 Job Scheduler O&M overview.....	644
3.1.5.3.2.2 Job Scheduler overview.....	645
3.1.5.3.2.3 Job Scheduler health.....	646
3.1.5.3.2.4 Job Scheduler quota management.....	647
3.1.5.3.2.5 Job Scheduler instances.....	649
3.1.5.3.2.6 Job Scheduler compute nodes.....	649
3.1.5.3.2.7 Enable or disable SQL acceleration.....	650
3.1.5.3.2.8 Restart the primary master node of Job Scheduler.....	651
3.1.5.3.3 Apsara Distribute File System O&M.....	652
3.1.5.3.3.1 Apsara Distribute File System O&M overview.....	652
3.1.5.3.3.2 Apsara Distributed File System overview.....	653
3.1.5.3.3.3 Service instances.....	655
3.1.5.3.3.4 Service health.....	655
3.1.5.3.3.5 Apsara Distributed File System storage nodes.....	656
3.1.5.3.3.6 Change the primary master node for Apsara Distributed File System.....	657
3.1.5.3.3.7 Empty the recycle bin of Apsara Distributed File System....	658
3.1.5.3.3.8 Enable or disable data rebalancing for Apsara Distributed File System.....	659
3.1.5.3.3.9 Run a checkpoint on the master nodes of Apsara Distributed File System.....	661
3.1.5.3.4 DataWorks O&M.....	661
3.1.5.3.4.1 DataWorks O&M overview.....	662
3.1.5.3.4.2 DataWorks overview.....	662
3.1.5.3.4.3 DataWorks health.....	663
3.1.5.3.4.4 DataWorks instances.....	664
3.1.5.3.4.5 DataWorks slots.....	664
3.1.5.3.4.6 DataWorks tasks.....	667
3.1.5.3.4.7 DataWorks configuration.....	668
3.1.5.3.4.8 DataWorks cluster scaling.....	668
3.1.5.3.5 Tunnel service.....	671
3.1.5.3.5.1 Tunnel service O&M overview.....	672
3.1.5.3.5.2 Tunnel service overview.....	673
3.1.5.3.5.3 Tunnel service instances.....	674
3.1.5.3.5.4 Restart tunnel servers.....	674

3.1.5.4 Cluster O&M.....	676
3.1.5.4.1 Cluster O&M overview.....	676
3.1.5.4.2 Cluster overview.....	677
3.1.5.4.3 Cluster health.....	683
3.1.5.4.4 Cluster hosts.....	688
3.1.5.4.5 Cluster scaling.....	689
3.1.5.5 Host O&M.....	693
3.1.5.5.1 Host O&M overview.....	693
3.1.5.5.2 Host overview.....	694
3.1.5.5.3 Host charts.....	700
3.1.5.5.4 Host health.....	700
3.1.5.5.5 Host services.....	705
3.1.6 StreamCompute.....	705
3.1.6.1 O&M overview.....	705
3.1.6.2 Business O&M.....	707
3.1.6.2.1 Projects.....	707
3.1.6.2.2 Jobs.....	707
3.1.6.2.3 Queues.....	708
3.1.6.3 Service O&M.....	709
3.1.6.3.1 Yarn.....	709
3.1.6.3.2 HDFS.....	710
3.1.6.4 Cluster O&M.....	711
3.1.6.4.1 Cluster overview.....	711
3.1.6.4.2 Cluster health.....	716
3.1.6.4.3 Hosts.....	721
3.1.6.4.4 Cluster scale-out.....	721
3.1.6.4.5 Cluster scale-in.....	723
3.1.6.5 Host O&M.....	725
3.1.6.5.1 Host overview.....	725
3.1.6.5.2 Host health.....	731
3.1.6.5.3 Host charts.....	735
3.1.6.5.4 Host services.....	736
3.1.7 DataHub.....	736
3.1.7.1 O&M overview.....	737
3.1.7.2 Business O&M.....	740
3.1.7.2.1 Business O&M overview.....	740
3.1.7.2.2 Projects.....	741
3.1.7.2.3 Topics.....	743
3.1.7.3 Service O&M.....	745
3.1.7.3.1 Service O&M for Job Scheduler.....	745
3.1.7.3.1.1 Service O&M overview.....	745
3.1.7.3.1.2 Service overview.....	745
3.1.7.3.1.3 Service instances.....	747
3.1.7.3.1.4 Health status.....	747
3.1.7.3.1.5 Compute nodes.....	748

3.1.7.3.2 Service O&M for Apsara Distributed File System.....	749
3.1.7.3.2.1 Service O&M overview.....	749
3.1.7.3.2.2 Service overview.....	750
3.1.7.3.2.3 Service instances.....	751
3.1.7.3.2.4 Health status.....	752
3.1.7.3.2.5 Storage nodes.....	752
3.1.7.3.2.6 Empty the recycle bin of Apsara Distributed File System....	753
3.1.7.3.2.7 Enable or disable data rebalancing for Apsara Distributed File System.....	754
3.1.7.3.2.8 Run a checkpoint on the master nodes of Apsara Distributed File System.....	755
3.1.7.3.2.9 Change the primary master node for Apsara Distributed File System.....	756
3.1.7.4 Cluster O&M.....	758
3.1.7.4.1 Cluster O&M overview.....	758
3.1.7.4.2 Cluster overview.....	759
3.1.7.4.3 Cluster health.....	763
3.1.7.4.4 Cluster hosts.....	768
3.1.7.4.5 Cluster scale-out.....	769
3.1.7.4.6 Cluster scale-in.....	771
3.1.7.4.7 Delete topics from a smoke testing project.....	773
3.1.7.4.8 Reverse parse RequestId.....	774
3.1.7.5 Host O&M.....	775
3.1.7.5.1 Host O&M overview.....	775
3.1.7.5.2 Host overview.....	776
3.1.7.5.3 Host charts.....	781
3.1.7.5.4 Host health.....	781
3.1.7.5.5 Host services.....	785
3.1.8 Elasticsearch.....	786
3.1.8.1 O&M overview.....	786
3.1.8.2 Business O&M.....	787
3.1.8.2.1 Cluster configuration.....	787
3.1.8.2.2 System configuration.....	789
3.1.8.3 Service O&M.....	789
3.1.8.3.1 Service overview.....	789
3.1.8.3.2 Service hosts.....	793
3.1.8.4 Cluster O&M.....	793
3.1.8.4.1 Cluster overview.....	793
3.1.8.4.2 Cluster health.....	799
3.1.8.5 Host O&M.....	804
3.1.8.5.1 Host overview.....	804
3.1.8.5.2 Host charts.....	809
3.1.8.5.3 Host health.....	809
3.1.8.5.4 Host services.....	813
3.1.9 Dataphin.....	813

3.1.9.1 O&M overview.....	814
3.1.9.2 Service O&M.....	815
3.1.9.2.1 Service overview.....	815
3.1.9.2.2 Service hosts.....	820
3.1.9.3 Cluster O&M.....	820
3.1.9.3.1 Cluster overview.....	820
3.1.9.3.2 Cluster health.....	825
3.1.9.4 Host O&M.....	830
3.1.9.4.1 Host overview.....	830
3.1.9.4.2 Host health.....	836
3.1.10 I+.....	840
3.1.10.1 O&M overview.....	840
3.1.10.2 Service O&M.....	842
3.1.10.2.1 Service overview.....	842
3.1.10.2.2 Service hosts.....	847
3.1.10.3 Cluster O&M.....	847
3.1.10.3.1 Cluster overview.....	847
3.1.10.3.2 Cluster health.....	852
3.1.10.4 Host O&M.....	857
3.1.10.4.1 Host overview.....	857
3.1.10.4.2 Host health.....	862
3.1.11 Quick BI.....	867
3.1.11.1 O&M overview.....	867
3.1.11.2 Service O&M.....	868
3.1.11.2.1 Service overview.....	869
3.1.11.2.2 Service hosts.....	873
3.1.11.3 Cluster O&M.....	874
3.1.11.3.1 Cluster overview.....	874
3.1.11.3.2 Cluster health.....	878
3.1.11.4 Host O&M.....	883
3.1.11.4.1 Host overview.....	883
3.1.11.4.2 Host health.....	889
3.1.12 PAI.....	893
3.1.12.1 O&M overview.....	893
3.1.12.2 Service O&M.....	895
3.1.12.2.1 Service overview.....	895
3.1.12.2.2 Service hosts.....	900
3.1.12.3 Cluster O&M.....	900
3.1.12.3.1 Cluster overview.....	901
3.1.12.3.2 Cluster health.....	905
3.1.12.4 Host O&M.....	910
3.1.12.4.1 Host overview.....	910
3.1.12.4.2 Host health.....	916
3.1.13 Management.....	920
3.1.13.1 Overview.....	920

3.1.13.2 Jobs.....	921
3.1.13.2.1 Overview.....	921
3.1.13.2.2 Jobs.....	923
3.1.13.2.2.1 Run a job from a scheme.....	923
3.1.13.2.2.2 Create a job from a scheme.....	926
3.1.13.2.2.3 Enable or disable a cron job.....	933
3.1.13.2.2.4 Manually run a job.....	934
3.1.13.2.2.5 View jobs.....	936
3.1.13.2.2.6 View the execution history of a job.....	937
3.1.13.2.3 Schemes.....	938
3.1.13.2.3.1 Create a scheme from a job.....	938
3.1.13.2.3.2 View schemes.....	939
3.1.13.2.3.3 View the execution history of a scheme.....	940
3.1.13.2.4 View the execution history.....	941
3.1.13.3 Patch management.....	946
3.1.13.4 Hot upgrade.....	948
3.1.13.5 Health management.....	950
3.1.13.6 Operation auditing.....	954
3.1.14 Go to other platforms.....	956
3.2 MaxCompute.....	957
3.2.1 Concepts and architecture.....	957
3.2.2 O&M commands and tools.....	961
3.2.2.1 Before you start.....	961
3.2.2.2 odpscmd commands.....	961
3.2.2.3 Tunnel commands.....	966
3.2.2.4 LogView tool.....	972
3.2.2.4.1 Before you start.....	972
3.2.2.4.2 LogView introduction.....	975
3.2.2.4.3 Preliminary knowledge of LogView.....	975
3.2.2.4.4 Basic operations and examples.....	981
3.2.2.4.5 Best practices.....	984
3.2.2.5 Apsara Bigdata Manager.....	984
3.2.3 Routine O&M.....	986
3.2.3.1 Configurations.....	986
3.2.3.2 Routine inspections.....	987
3.2.3.3 Chunkserver and tunnel scale-in.....	993
3.2.3.4 Chunkserver and tunnel scale-out.....	1000
3.2.3.5 Shut down a chunkserver, perform maintenance, and then clone the chunkserver.....	1011
3.2.3.6 Shut down a chunkserver for maintenance without compromising the system.....	1017
3.2.3.7 Adjust the virtual resources of the Apsara system in MaxCompute.....	1018
3.2.3.8 Restart a MaxCompute service.....	1023
3.2.4 Common issues and solutions.....	1026

3.2.4.1	View and allocate MaxCompute cluster resources.....	1026
3.2.4.2	Common issues and data skew troubleshooting.....	1039
3.3	DataWorks.....	1049
3.3.1	Basic concepts and structure.....	1050
3.3.1.1	What is DataWorks (base)?.....	1050
3.3.1.2	Functions of base.....	1050
3.3.1.3	Introduction to data analytics.....	1050
3.3.1.4	Architecture of DataWorks in Apsara Stack V3.....	1052
3.3.1.5	Directory of each service.....	1054
3.3.2	Common administration tools and commands.....	1055
3.3.2.1	Find the container that runs the service.....	1055
3.3.2.2	Cluster resource list.....	1055
3.3.2.3	Commands to restart services.....	1056
3.3.2.4	View the log of a failed task.....	1056
3.3.2.5	Rerun a task.....	1056
3.3.2.6	Terminate a task.....	1056
3.3.2.7	Filter tasks in the administration center.....	1057
3.3.2.8	Commonly used Linux commands.....	1057
3.3.2.9	View the slots usage of each resource group.....	1058
3.3.3	Process daily administration operations.....	1059
3.3.3.1	Daily check.....	1059
3.3.3.1.1	Check the service status and the basic information of the servers.....	1059
3.3.3.1.2	Check the postgres database.....	1060
3.3.3.1.3	Check the status of each gateway server.....	1060
3.3.3.1.4	Check the case test report.....	1061
3.3.3.2	View logs of the services.....	1061
3.3.3.3	Scale out the node cluster that runs the base-biz-gateway service.....	1061
3.3.3.4	Scale in the node cluster that runs the base-biz-gateway service.....	1066
3.3.3.5	Restart the base-biz-alisa service.....	1069
3.3.3.6	Restart the base-biz-phoenix service.....	1070
3.3.3.7	Restart base-biz-tenant.....	1070
3.3.3.8	Restart base-biz-gateway.....	1071
3.3.3.9	Restart the base-biz-api service.....	1072
3.3.3.10	Restart the base-redis service.....	1072
3.3.3.11	Restart DataWorks Data Service.....	1073
3.3.3.12	Restart DataWorks Data Management.....	1074
3.3.4	Common issues and solutions.....	1074
3.3.4.1	Node instances remain in the Pending (Resources) status.....	1074
3.3.4.2	An out-of-memory (OOM) error occurs when synchronizing data from an Oracle database.....	1078
3.3.4.3	A task does not run at the specified time.....	1078
3.3.4.4	The test service of base is not in the desired status.....	1079

3.3.4.5 The Data Management page does not display the number of tables and the usage of tables.....	1079
3.3.4.6 Logs are not automatically cleaned up.....	1080
3.3.4.7 The real-time analysis service is not in the desired status.....	1081
3.4 Realtime Compute.....	1081
3.4.1 Job status.....	1081
3.4.1.1 Overview.....	1081
3.4.1.2 Task status.....	1082
3.4.1.3 Health score.....	1082
3.4.1.4 Job instantaneous values.....	1082
3.4.1.5 Running topology.....	1083
3.4.2 Curve charts.....	1086
3.4.2.1 Overview.....	1086
3.4.2.2 Overview.....	1087
3.4.2.3 Advanced view.....	1089
3.4.2.4 Processing delay.....	1092
3.4.2.5 Throughput.....	1092
3.4.2.6 Queue.....	1092
3.4.2.7 Tracing.....	1092
3.4.2.8 Process.....	1093
3.4.2.9 JVM.....	1093
3.4.3 FailOver.....	1094
3.4.4 CheckPoints.....	1094
3.4.5 JobManager.....	1095
3.4.6 TaskExecutor.....	1095
3.4.7 Data lineage.....	1095
3.4.8 Properties and Parameters.....	1096
3.4.9 Improve performance by automatic configuration.....	1097
3.4.10 Improve performance by manual configuration.....	1104
3.4.10.1 Overview.....	1104
3.4.10.2 Optimize resource configuration.....	1104
3.4.10.3 Improve performance based on job parameter settings.....	1107
3.4.10.4 Optimize upstream and downstream data storage based on parameter settings.....	1107
3.4.10.5 Apply new configuration.....	1108
3.4.10.6 Concepts.....	1109
3.5 Apsara Bigdata Manager (ABM).....	1110
3.5.1 Routine maintenance.....	1110
3.5.1.1 Perform routine maintenance.....	1110
3.5.1.2 View the ABM operating status.....	1111
3.5.1.3 Troubleshooting.....	1115
3.5.2 Backup and restore.....	1115
3.6 Quick BI.....	1116
3.6.1 Introduction to O&M and tools.....	1116
3.6.1.1 Introduction to operations and maintenance.....	1116

3.6.1.2	Troubleshoot Quick BI issues by using the Apsara Infrastructure Management Framework.....	1116
3.6.2	Routine maintenance.....	1119
3.6.2.1	Introduction to Quick BI components.....	1119
3.6.2.2	Database initialization components.....	1120
3.6.2.3	Cache components.....	1121
3.6.2.4	Runtime components.....	1121
3.6.2.5	Web service components.....	1122
3.6.2.6	Automated testing components.....	1123
3.7	Graph Analytics.....	1124
3.7.1	Operations and maintenance tools and logon methods.....	1124
3.7.1.1	Log on to Apsara BigData Manager.....	1124
3.7.1.2	Log on to Apsara Infrastructure Management Framework.....	1126
3.7.1.3	Log on to the Graph Analytics container.....	1128
3.7.2	Operations and maintenance.....	1130
3.7.2.1	Operations and maintenance based on BigData Manager.....	1130
3.7.2.1.1	View and handle cluster alerts.....	1130
3.7.2.1.2	View cluster performance metrics.....	1134
3.7.2.1.3	View server operation metrics.....	1135
3.7.2.2	Operations and maintenance based on Apsara Infrastructure Management Framework.....	1136
3.7.2.3	Operations and maintenance based on the Graph Analytics container.....	1139
3.7.2.3.1	View instances.....	1139
3.7.2.3.2	Log files.....	1140
3.7.2.3.3	Database logs.....	1140
3.7.2.3.4	Stop the service.....	1141
3.7.2.3.5	Restart the service.....	1142
3.7.3	Security maintenance.....	1142
3.7.3.1	Network security maintenance.....	1142
3.7.3.2	Account password maintenance.....	1142
3.7.4	Troubleshooting.....	1143
3.7.4.1	Fault response mechanism.....	1143
3.7.4.2	Troubleshooting methods.....	1143
3.7.4.3	Common failure troubleshooting.....	1143
3.7.4.4	Hardware troubleshooting.....	1143
3.8	Machine Learning Platform for AI.....	1144
3.8.1	Query server and application information.....	1144
3.8.1.1	Apsara Stack Machine Learning Platform for AI.....	1144
3.8.1.1.1	Query server information.....	1144
3.8.1.1.2	Log on to a server.....	1145
3.8.1.1.3	Query configurations.....	1145
3.8.1.1.4	Restart an application service.....	1146
3.8.1.2	Online model service.....	1147
3.8.1.2.1	Query online model service information.....	1147

3.8.1.2.2 Log on to the online model service container.....	1147
3.8.1.2.3 Restart a pod.....	1148
3.8.1.3 GPU cluster and task information.....	1148
3.8.1.3.1 Query GPU cluster information.....	1148
3.8.1.3.2 Query GPU task information.....	1148
3.8.2 Maintenance and troubleshooting.....	1149
3.8.2.1 Machine Learning Platform for AI maintenance.....	1149
3.8.2.1.1 Run ServiceTest.....	1149
3.8.2.1.2 Common faults and solutions.....	1150
3.8.2.1.2.1 Maintenance commands.....	1150
3.8.2.1.2.2 pai.xx.xx access failures.....	1150
3.8.2.1.2.3 Experiment failures.....	1152
3.8.2.1.2.4 Other failures.....	1153
3.8.2.2 Online model service maintenance (must be activated separately).....	1153
3.8.2.3 GPU cluster maintenance (deep learning must be activated separately).....	1154
3.9 Elasticsearch.....	1155
3.9.1 Log on to the Apsara Stack Elasticsearch O&M center.....	1155
3.9.2 Apsara Stack Elasticsearch O&M center.....	1156
3.9.2.1 O&M overview.....	1156
3.9.2.2 Business O&M.....	1158
3.9.2.2.1 Cluster configuration.....	1158
3.9.2.2.2 System configuration.....	1159
3.9.2.3 Service O&M.....	1160
3.9.2.3.1 Service overview.....	1160
3.9.2.3.2 Service hosts.....	1164
3.9.2.4 Cluster O&M.....	1164
3.9.2.4.1 Cluster overview.....	1164
3.9.2.4.2 Cluster health.....	1169
3.9.2.5 Host O&M.....	1174
3.9.2.5.1 Host overview.....	1174
3.9.2.5.2 Host charts.....	1179
3.9.2.5.3 Host health.....	1180
3.9.2.5.4 Host services.....	1184
3.9.3 Online O&M.....	1185
3.9.3.1 Cluster health.....	1185
3.9.4 Troubleshooting.....	1185
3.9.4.1 The cluster health status is yellow.....	1185
3.9.4.2 Query index status.....	1186
3.9.4.3 Restore index status.....	1186
3.10 DataHub.....	1186
3.10.1 Concepts and architecture.....	1186
3.10.1.1 Terms.....	1186
3.10.1.2 Architecture.....	1190

3.10.1.2.1 Feature oriented architecture.....	1190
3.10.1.2.2 Technical architecture.....	1192
3.10.2 Commands and tools.....	1193
3.10.2.1 Common commands for the Apsara system.....	1193
3.10.2.2 Common commands for Apsara Distributed File System.....	1194
3.10.2.3 Common commands for Job Scheduler.....	1196
3.10.2.4 Xstream.....	1197
3.10.3 Routine maintenance.....	1201
3.10.3.1 Remove chunk servers from the DataHub cluster.....	1201
3.10.3.2 Add chunk servers to the DataHub cluster.....	1203
3.10.3.3 Restore data after a power outage.....	1207
3.10.3.4 Reimage frontend and chunk servers.....	1208
3.10.3.5 Upgrade or redeploy DataHub.....	1209
3.10.3.6 Shut down problematic chunk servers.....	1210
3.10.3.7 Shut down the DataHub cluster.....	1214
3.10.3.8 Replace a hard drive with a new one on the pangu_cs node..	1215
3.10.4 Exceptions and solutions.....	1216
3.10.5 Appendix.....	1218
3.10.5.1 Installation environment.....	1218
3.10.5.2 Deployment directories and services.....	1218
3.10.5.3 Error codes.....	1220

1 Operations of basic platforms

1.1 Apsara Stack Operations (ASO)

1.1.1 Apsara Stack Operations overview

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

ASO has the following main functions:

- **Operations and Maintenance Dashboard**

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

- **Alarm Monitoring**

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

- **Resource Management**

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

- **Inventory Management**

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

- **Products**

Products allows you to access the operations and maintenance services of other products on the cloud platform. You are redirected to the corresponding

operations and maintenance page of a product by using Single Sign-On (SSO) and redirection.

- **ITIL Management**

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

- **Configurations**

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

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

- **Offline Backup**

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

- **NOC**

Network Operation Center (NOC) provides the operations capabilities such as the visualization of network-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.

- **Full Stack Monitor**

Full Stack Monitor 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.

- **Pangu Operation Center**

Pangu Operation Center displays the pangu graph, cluster information, node information, and pangu cluster status.

- **Task Management**

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

- **System Management**

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

1.1.2 Log on to Apsara Stack Operations

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

Prerequisites

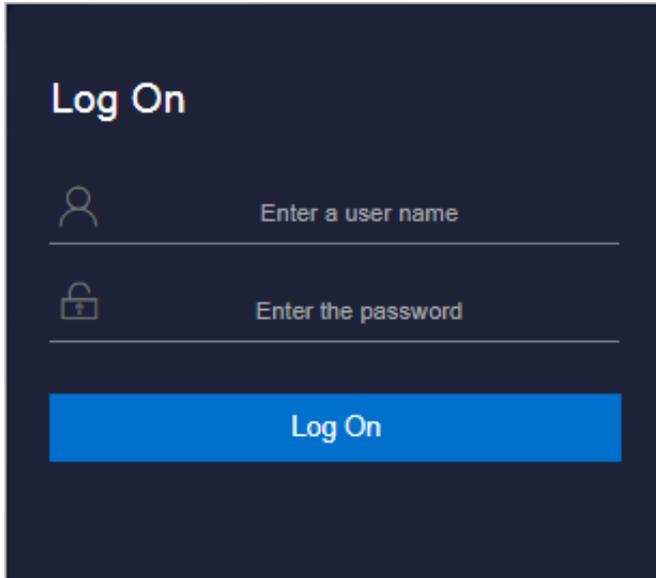
- **ASO access address in the format of `http://region-id.aso.intranet-domain-id.com`.**
- **Google Chrome browser (recommended).**

Procedure

- 1. Open the browser.**

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

Figure 1-1: Log on to ASO



Note:

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

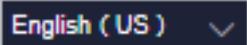
3. Enter the correct username and password.
 - The system has three default users:
 - Security officer: manages other users or roles.
 - Auditor officer: views audit logs.
 - System administrator: used for other functions except those of the security officer and auditor officer.
 - You must modify the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

1.1.3 Web page introduction

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



The description of each area is as follows.

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

1.1.4 Operation and maintenance dashboard

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

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

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

1.1.5 Alarm Monitoring

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

1.1.5.1 Overview

Alarm Monitoring allows you to view the overview information of alerts.

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, select **Alarm Monitoring**.
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, enter the filter name and then configure the filter conditions.

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

- **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.

1.1.5.2 Alert events

Alert Events 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 Operation & Maintenance (O&M) operations on the alerts.

Context

Alert Events 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, NAS, MaxCompute, DataHub, Graph Analytics, yundun-advance, yundun-common, EDAS, QuickBI, Elasticsearch, and Table Store.
- **Timeout Alert:** Displays the information of all the timeout alerts.

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. In the left-side navigation pane, choose **Alarm Monitoring > Alert Events.**

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/or 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.

- 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.

- 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 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.

- Shield alerts

For alerts that do not need to be focused on, you can select them and then click Shield at the top of the page to shield the alerts. Then, such alerts are not displayed in the system.

- Remove the shield

To remove the shield after you shield the alerts, select the shielded alerts and then click Remove Shield at the top of the page.

- Export a report

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

1.1.5.3 Alert history

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Alarm 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/or search content.**

- **Export a list of alerts**

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

- **View alert sources**

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

- **View alert details**

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

- **View the original alert information**

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

1.1.5.4 Alert configuration

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

1.1.5.4.1 Alert contacts

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

Procedure

1. *Log on to Apsara Stack Operations.*

2. In the left-side navigation pane, choose **Alarm 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/or phone number and then click **Search**. The alert contacts that meet the search conditions are displayed in the list.

- **Add an alert contact**

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

- **Modify an alert contact**

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

- **Delete an alert contact**

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

1.1.5.4.2 Alert contact groups

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

Procedure

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

4. Then, you can:

- **Search for an alert contact group**

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

- **Add an alert contact group**

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

- **Modify an alert contact group**

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

- **Delete one or more alert contact groups**

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

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

1.1.5.4.3 Static parameter settings

You can configure the static parameters related to alerts based on business needs.

Currently, you can only configure the parameter related to timeout alerts.

Context

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose Alarm 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	<p>The default value is 5, indicating 5 days.</p> <p>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 Alarm Monitoring > Alert Events.</p> <p>For example, if the parameter value is 5, the system displays the alert events that exceed 5 days on the Timeout Alert tab of Alarm Monitoring > Alert Events.</p>
Description	Enter the description related to the configuration.

7. Then, click OK.

1.1.5.5 Alert overview

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

Procedure

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

1.1.5.6 Alert subscription and push

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

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Alarm Monitoring > Subscribe/Push**.
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. <ul style="list-style-type: none"> • 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.

Configuration	Description
Push Mode	<p>Select one of the following methods:</p> <ul style="list-style-type: none"> • ALL: All of the alerts are pushed in each push cycle. • TOP: Only alerts with high priority are pushed in each push cycle.
Push Template	<p>Select one of the following templates:</p> <ul style="list-style-type: none"> • 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. <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: A preset ANS template exists if the system already connects with the ANS product. To restore the initial configurations of the template with one click, click Reset. </div>
Custom JSON Fields	<p>The person who receives the push can use this field to configure the identifier in a custom way. The format must be JSON.</p>
Push Switch	<p>Select whether to push the alerts.</p> <p>If the switch is not turned on here, you can enable the push feature in the Push Switch column after configuring the subscription channel.</p>

5. After completing the configurations, click OK.

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

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



Note:

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

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

1.1.6 Resource Management

Resource Management monitors and manages physical servers in the data center. The major monitoring information includes the overall status information, monitoring metrics, alert delivery status, and port traffic of devices.

1.1.6.1 Physical servers

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

Procedure

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

In the upper-right corner of the page, you can view the number of physical servers, the number of servers with alerts, and the number of alerts.



Note:

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

3. Click the Product tab and then you can:

- **Expand the left-side navigation tree level by level based on regions, products, and clusters to view a list of physical servers where a service of a product is located on the right.**
- **You can search for and view a physical server by product, cluster, group, or hostname in the search box on the right.**
- **Click Details in the Operation column at the right of a product to go to the Physical Server Details page, and then view the basic information, monitoring details, and alarm information of the physical server.**

You can view the monitoring details or alarm information by switching between the two tabs, and select different time ranges to view the monitoring metrics. The monitoring metrics include CPU usage, system load, disk usage, memory usage, network throughput, and disk I/O.

- **Click  in the upper-right corner to export the information of all physical servers to your local computer.**

4. Click the Server tab and then you can:

- **Expand the left-side navigation tree level by level based on data centers and racks to view a list of physical servers in a rack on the right.**
- **You can search for and view a physical server by hostname, IP address, device function, or SN in the search box on the right.**
- **Click  in the upper-right corner and then enter the physical server information to add a physical server.**
- **Click Details in the Operation column at the right of a physical server to go to the Physical Server Details page, and then view the basic information, monitoring details, and alarm information of the physical server.**

You can view the monitoring details or alarm information by switching between the two tabs, and select different time ranges to view the monitoring

metrics. The monitoring metrics include CPU usage, system load, disk usage, memory usage, network throughput, and disk I/O.

- Click **Modify** in the Operation column at the right of a physical server to update the physical server information.
- Click **Delete** in the Operation column at the right of a physical server to delete the physical server if it is not required to be monitored.
- Click  in the upper-right corner to export the information of all physical servers to your local computer.

1.1.7 Inventory Management

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

1.1.7.1 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 Operation & Maintenance (O&M) operations according to actual requirements.

Procedure

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



Note:

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

3. View the ECS inventory.
 - **CPU Inventory Details(Core)** and **Memory Inventory Details(TB)** display the used and available CPU (core) and memory (TB) of all ECS instance type families in the last five days.
 - **ECS Instances Inventory Details** allows you to perform a paging query on the inventory details of a certain type of ECS instances at a certain date by Zone, Instance Type, and Date. For more information about the mapping between

instance type families and CPU/memory configurations of instances, see [Table 1-1: Instance type](#).

Table 1-1: Instance type

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

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

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
d1	ecs.d1.2xlarge	4 * 5500	8	32.0	3
	ecs.d1.4xlarge	8 * 5500	16	64.0	8
	ecs.d1.6xlarge	12 * 5500	24	96.0	8
	ecs.d1-c8d3.8xlarge	12 * 5500	32	128.0	8
	ecs.d1.8xlarge	16 * 5500	32	128.0	8
	ecs.d1-c14d3.14xlarge	12 * 5500	56	160.0	8
	ecs.d1.14xlarge	28 * 5500	56	224.0	8
gn4	ecs.gn4-c4g1.xlarge	None	4	30.0	2
	ecs.gn4-c8g1.2xlarge	None	8	60.0	3
	ecs.gn4.8xlarge	None	32	48.0	8
	ecs.gn4-c4g1.2xlarge	None	8	60.0	3
	ecs.gn4-c8g1.4xlarge	None	16	60.0	8
	ecs.gn4.14xlarge	None	56	96.0	8

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
ga1	ecs.ga1.xlarge	1*87	4	10.0	2
	ecs.ga1.2xlarge	1*175	8	20.0	3
	ecs.ga1.4xlarge	1*350	16	40.0	8
	ecs.ga1.8xlarge	1*700	32	80.0	8
	ecs.ga1.14xlarge	1*1400	56	160.0	8
se1ne	ecs.se1ne.large	None	2	16.0	1
	ecs.se1ne.xlarge	None	4	32.0	2
	ecs.se1ne.2xlarge	None	8	64.0	3
	ecs.se1ne.4xlarge	None	16	128.0	8
	ecs.se1ne.8xlarge	None	32	256.0	8
	ecs.se1ne.14xlarge	None	56	480.0	8
sn2ne	ecs.sn2ne.large	None	2	8.0	1
	ecs.sn2ne.xlarge	None	4	16.0	2
	ecs.sn2ne.2xlarge	None	8	32.0	3
	ecs.sn2ne.4xlarge	None	16	64.0	8

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	ecs.sn2ne.8xlarge	None	32	128.0	8
	ecs.sn2ne.14xlarge	None	56	224.0	8
sn1ne	ecs.sn1ne.large	None	2	4.0	1
	ecs.sn1ne.xlarge	None	4	8.0	2
	ecs.sn1ne.2xlarge	None	8	16.0	3
	ecs.sn1ne.4xlarge	None	16	32.0	8
	ecs.sn1ne.8xlarge	None	32	64.0	8
gn5i	ecs.gn5i-c2g1.large	None	2	8.0	1
	ecs.gn5i-c4g1.xlarge	None	4	16.0	2
	ecs.gn5i-c8g1.2xlarge	None	8	32.0	2
	ecs.gn5i-c16g1.4xlarge	None	16	64.0	2
	ecs.gn5i-c28g1.14xlarge	None	56	224.0	2
g5	ecs.g5.large	None	2	8.0	2
	ecs.g5.xlarge	None	4	16.0	3

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	ecs.g5.2xlarge	None	8	32.0	4
	ecs.g5.4xlarge	None	16	64.0	8
	ecs.g5.6xlarge	None	24	96.0	8
	ecs.g5.8xlarge	None	32	128.0	8
	ecs.g5.16xlarge	None	64	256.0	8
	ecs.g5.22xlarge	None	88	352.0	15
c5	ecs.c5.large	None	2	4.0	2
	ecs.c5.xlarge	None	4	8.0	3
	ecs.c5.2xlarge	None	8	16.0	4
	ecs.c5.4xlarge	None	16	32.0	8
	ecs.c5.6xlarge	None	24	48.0	8
	ecs.c5.8xlarge	None	32	64.0	8
	ecs.c5.16xlarge	None	64	128.0	8
r5	ecs.r5.large	None	2	16.0	2
	ecs.r5.xlarge	None	4	32.0	3
	ecs.r5.2xlarge	None	8	64.0	4

Instance type family	Type code	Local storage (GiB)	CPU (core)	Memory (GiB)	Elastic Network Interface (ENI) (including a primary ENI)
	ecs.r5.4xlarge	None	16	128.0	8
	ecs.r5.6xlarge	None	24	192.0	8
	ecs.r5.8xlarge	None	32	256.0	8
	ecs.r5.16xlarge	None	64	512.0	8
	ecs.r5.22xlarge	None	88	704.0	15
se1	ecs.se1.large	None	2	16.0	2
	ecs.se1.xlarge	None	4	32.0	3
	ecs.se1.2xlarge	None	8	64.0	4
	ecs.se1.4xlarge	None	16	128.0	8
	ecs.se1.8xlarge	None	32	256.0	8
	ecs.se1.14xlarge	None	56	480.0	8
d1ne	ecs.d1ne.2xlarge	4 * 5500	8	32.0	4
	ecs.d1ne.4xlarge	8 * 5500	16	64.0	8
	ecs.d1ne.6xlarge	12 * 5500	24	96.0	8
	ecs.d1ne.8xlarge	16 * 5500	32	128.0	8

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

1.1.7.2 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 Operation & Maintenance (O&M) operations according to actual requirements.

Procedure

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



Note:

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

3. **View the SLB inventory.**
 - The section in the upper-left corner displays the frozen, assigned, protected, and released internal VIP history inventory and public VIP history inventory in the last five days.
 - The section in the upper-right corner displays the used inventory and the corresponding percentage of the internal VIP and the public VIP.
 - SLB Inventory Details allows you to perform a paging query on the SLB inventory details by Type and Date.

1.1.7.3 View the RDS inventory

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

Procedure

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



Note:

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

3. View the RDS inventory.

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

1.1.7.4 View the OSS inventory

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

Procedure

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



Note:

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

3. View the OSS inventory.

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

1.1.7.5 View the Table Store inventory

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

Procedure

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



Note:

You can click  in the upper-right corner to configure the global quota.

3. View the Table Store inventory.

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

1.1.7.6 View the Log Service inventory

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

Procedure

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



Note:

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

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

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

1.1.7.7 View the EBS inventory

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Inventory Management > EBS**.
3. If multiple ECS clusters exist in the environment, click the tab of the corresponding ECS cluster to view the EBS inventory.
 - **Inventory Availability History(TB)** displays the available EBS inventory in the last five days.
 - **Current Inventory Usage(TB)** displays the used EBS inventory and the corresponding percentage.
 - **EBS Inventory Details** allows you to perform a paging query on the EBS inventory details by Date.

1.1.7.8 View the NAS inventory

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

Procedure

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

3. View the NAS inventory.

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

1.1.7.9 View the HDFS inventory

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Inventory Management > DFS**.
3. View the HDFS inventory.
 - **Inventory Availability History(TB)** displays the available HDFS inventory in the last five days.
 - **Current Inventory Usage(TB)** displays the used HDFS inventory and the corresponding percentage.
 - **DFS Inventory Details** allows you to perform a paging query on the HDFS inventory details by Date.

1.1.8 Products

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

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

In the product list, you can view the operations and maintenance icons of different products based on your permissions. For example, a Table Store operations engineer can only view the **OTS Storage Operations and Maintenance System** icon. Click **OTS Storage Operations and Maintenance System** to go to the Table Store operations and maintenance console. An operations system administrator can view all the operations and maintenance components of the cloud platform. The read

and write permissions for product operations and maintenance are separated. Therefore, different permissions can be dynamically assigned based on different roles.

1.1.9 ITIL Management

1.1.9.1 Overview

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

ITIL has the following functions:

- **Dashboard**

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

- **Services**

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

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

- **Incident management:** used to recover from exceptions and guarantee the normal production by a series of recovery operations, including diagnosis, processing, resolution, and confirmation. Incident management provides a unified mode and standardizes the process for incident processing, and supports automatically collecting or manually recording the incident information.
- **Problem management:** Temporarily resolved incidents or incidents whose root cause is not clear can be transformed to problems for further analysis and thorough troubleshooting. Incidents aim to resume the production, whereas problems aim to be completely solved to make sure the problems do

not recur. Problem management allows you to find the root cause of incidents, thoroughly troubleshoot the incidents, and reduce repeated incidents.

- **Version control**

Version Control displays the version information of Apsara Stack products.

- **Process template configuration**

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

- **CAB/ECAB configuration**

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

1.1.9.2 Dashboard

Dashboard allows you to view the summary of incident requests, problem requests, and change requests, namely the total numbers of incident requests, problem requests, and change requests, the numbers of new and closed incident requests, problem requests, and change requests, and their change trend. You can also view the distribution of request fulfillment and the information of version management.

Procedure

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

1.1.9.3 Services

1.1.9.3.1 Basic functions

1.1.9.3.1.1 Overview

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

Services is composed of requests and tasks.

- **Requests**

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

- **Tasks**

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

1.1.9.3.1.2 Manage requests

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

Procedure

1. *Log on to Apsara Stack Operations.*

2. In the left-side navigation pane, choose **ITIL Management > Services**. Click the **Request** tab.

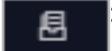
3. On the **Request** tab, you can:

- **Create a request**

Click  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.

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

- Filter requests

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

- Search for requests

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

- View request details

Find the request that you are about to view the details, and then click Detail.

The request details page is composed of the following sections:

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

1.1.9.3.1.3 Manage tasks

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

Context

Tasks are currently divided into the following three types:

- : My task, indicating tasks that are waiting to be processed by you.
- : Task pool, indicating a collection of tasks that are not assigned to related person in charge. You can check out the tasks in the task pool to make the tasks exclusive to you. Others cannot process the tasks that you have checked out. You can view the checked out tasks under .
- : 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 Apsara Stack Operations.](#)
2. In the left-side navigation pane, choose ITIL Management > Services. Click the My Task tab.
3. On the My Task tab, you can:
 - Search for tasks

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

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

1.1.9.3.2 Manage incidents

1.1.9.3.2.1 Create an incident request

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

Context

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

- **Automatically created**

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

- **Manually created**

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

Procedure

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

Configuration	Description
Priority	<p>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:</p> <ul style="list-style-type: none"> • Critical • Major • Minor • Remind • Cleared • System
Alarm Code	The alert ID.
Summary	The summary of this request.
Description	The detailed description about the request.
Suggestion	Optional. The suggestion about the request processing.

4. After completing the configurations, click Confirm.

1.1.9.3.2 Manage incident requests

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

Prerequisites

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

Procedure

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

5. On the request details page, you can:

- Change the priority

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



Note:

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

- Comment on the incident request

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

- Suspend the incident request

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

- Resume the incident request

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

- Recycle the incident request

Perform this operation for incident requests in the in processing () list.

Click Recycle to cancel or logically delete the incident request. The incident request is in the recycle bin () 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 () list.

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

1.1.9.3.2.3 Manage incident tasks

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

Context

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

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

Procedure

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

1.1.9.3.3 Manage problems

1.1.9.3.3.1 Create a problem request

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

Context

Temporarily resolved incidents or incidents whose root cause is not clear can be transformed to problems for further analysis and thorough troubleshooting. Problem management allows you to find the root 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 Apsara Stack Operations](#).
2. In the left-side navigation pane, choose ITIL Management > Services. Click the Request tab.

3. Click  and then select Problem. Configure the problem request on the 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: <ul style="list-style-type: none"> • Critical • Major • Minor • Remind • Cleared • System
Alarm Code	The alert ID.
Summary	The summary of this request.
Description	The detailed description about the request.
Suggestion	Optional. The suggestion about the request processing.

4. After completing the configurations, click Confirm.

1.1.9.3.3.2 Manage problem requests

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

Prerequisites

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

Procedure

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

- **Change the priority**

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



Note:

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

- **Comment on the problem request**

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

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

- **Suspend the problem request**

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

- **Resume the problem request**

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

- **Recycle the problem request**

Perform this operation for problem requests in the in processing () list. Click Recycle to cancel or logically delete the problem request. The problem request is in the recycle bin () list after being recycled.

- **Restore the problem request**

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

- **Delete the problem request**

Perform this operation for problem requests in the recycle bin () list. Click Delete to delete the problem request. After being deleted, the problem request is physically deleted and cannot be restored.

1.1.9.3.3.3 Manage problem tasks

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

Context

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

- **Diagnose:** analyzes the reason of the problem.

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

Procedure

1. *Log on to Apsara Stack Operations.*
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 problem has been processed after being reported because of the time difference. In this case, you can directly select Yes and configure the resolved date. Then, the Resolve phase is skipped and the system goes to the Confirm phase directly.

Configuration	Description
Remarks	The information about the task.

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

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

7. The system goes to the Confirm phase after the Resolve phase. This phase reviews the processing result of the problem. Then, click Confirm.
8. In the displayed Confirm dialog box, select the review result from the Is Pass drop-down list. Then, click OK.

The review results have the following three statuses:

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

1.1.9.4 Version control

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

Procedure

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



Note:

Before the search, click  to synchronize the information to Apsara Stack Operations (ASO).

1.1.9.5 Configure process templates

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

Log on to Apsara Stack Operations. In the left-side navigation pane, choose **ITIL Management > Process Template Configuration**. On this page, you can view the following three sections: **Process, Process Template, and Regulation**.

Process

Currently, the following processes are supported:

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

Process template

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

-  is the start node of the process. A process usually starts with the request creation.

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

The templates can be classified into the following three types:

- **Incidents and problems**
Incident and Problem. The whole process has the following phases: Record, Diagnose, Resolve, Confirm, and Close.
- **Request fulfillment**
Change Role, Create Identity, Reset Password, and Logout Identity. The whole process has the following phases: Record, Approve, Handle, and Close.
- **Change management**
Change, Version Upgrade, Hotfix Upgrade, and Configuration Upgrade. The whole process has the following phases: Record, Preliminary Approval, Information Modify, CAB Audit, ECAB Audit, Schedule Arrangement, Task Execution, Task Confirmation, Review, and Close.

Regulation

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

Currently, the system supports four regulations:

- Assign by role
- Assign by user

- **Assign by owner**
- **CAB/ECAB configuration**

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



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.**
- **If the selected role has only one user, only that user can view the current task in my task.**
- **If the selected role has more than one user, all the users under the selected role can view the current task in the task pool.**

- **Assign by user**

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

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

- **Assign by owner**

If Assign by Owner is selected, only the user who creates the process request can view the current task in my task. The person who creates the request is the owner of the request.

- **CAB/ECAB configuration**

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

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

1.1.9.6 Configure CAB or ECAB

The change management process has the CAB Audit and ECAB Audit phases.

Therefore, you must configure the CAB or ECAB.

Context

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

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

If no regulation is configured, all the users can generate the current task in my task by default. With one or more users configured, each configured user can generate the current task in my task, and the task can go to the next phase only after all the users configured in this phase finish the current task.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose ITIL Management > CAB/ECAB Configuration.
3. Click the CAB Configuration 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.

1.1.10 Configurations

1.1.10.1 Overview

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

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

1.1.10.2 Modify a configuration item of a product

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

Procedure

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

3. Enter the name of the product or configuration item in the Product or Configuration Name field. Click Search to check if the configuration item already exists in the list.

- The configuration item already exists in the list.

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

Click Modify in the Actions column. In the displayed Modify Configurations dialog box, modify the values and then click OK to modify the configuration item locally.

- The configuration item does not exist in the list.

You must add a configuration item as follows:

- a. Click Add in the upper-right corner.
- b. In the displayed Add Configuration dialog box, configure the information, such as Product, Configuration Name, Default Value, and Data Source Type, for the configuration item.
- c. Click OK.

Then, this configuration item is displayed in the list. You can search for and modify this configuration item.

4. After the configuration item is modified, click Apply in the Actions column to make the modifications take effect.

5. Optional: To import or export configuration items as a file, click Import or Export in the upper-right corner.



Note:

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

1.1.10.3 Restore the configuration value of a modified configuration item

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Configurations > Restore**.
3. On the Restore page, enter the name of the configuration item whose configuration value you want to roll back in the Configuration Name field and then click Search. All modification records of the configuration item appear in the list.
4. Find the record to be rolled back, and then click Restore in the Actions column.
5. Click OK in the displayed dialog box to restore the configuration value of the configuration item.

1.1.10.4 Manage kernel configurations

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Configurations > Kernel Configurations**.
3. On the Kernel Configurations page, you can:

- **Add a kernel configuration**

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

- **Modify a kernel configuration**

Find the kernel configuration to be modified. Click Modify in the Actions column. Modify the Kernel Configuration, Read Command, and Modify Command. Then, click Save.

- **Delete a kernel configuration**

Find the kernel configuration to be deleted. Click Delete in the Actions column. In the displayed dialog box, click OK.

1.1.10.5 Scan configurations

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

Prerequisites

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

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

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose Configurations > Kernel Configurations Actions.
3. On the Kernel Configurations Actions page, enter the hostname or IP address in the search box and then click Scan Configuration.
The scan results are displayed in the list.
4. Optional: To modify the scanned configuration value, click Modify to modify the Configuration Value. Click Save to modify the local value of the kernel configuration.
After the modification, click Apply to apply the local value of the kernel configuration to the corresponding host. To read the value of the kernel configuration on the host again, click Get.

1.1.11 Offline Backup

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

1.1.11.1 Service configuration

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

1.1.11.1.1 Configure the backup server

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

Procedure

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

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

Configuration	Description
Backup Server IP Address	<p>The IP address of the backup server.</p> <p>The backup server must meet the following requirements:</p> <ul style="list-style-type: none"> • 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	<p>The storage path of backup files on the backup server.</p> <p>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.</p>
Backup Retention	<p>The actual time (in days) that backup files are stored . The backup file that exceeds the time is to be deleted.</p>

1.1.11.1.2 Add a backup product

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

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. In the left-side navigation pane, select **Offline Backup**.
3. Choose **Service Configuration > Product Management**.
4. Click **Add** in the upper-right corner.

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

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

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

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

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

1.1.11.2 Backup service

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

1.1.11.2.1 Backup configuration

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

Prerequisites

Make sure that a product backup item is added. For more information about how to add a product backup item, see [Add a backup product](#).

Context

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

Procedure

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

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

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

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

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

1.1.11.2.2 View the backup details

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

Procedure

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

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

1.1.11.2.3 View the backup server status

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, select **Offline Backup**.
3. Choose **Backup Service > Service Status**.
4. On the **Service Status** page, view the memory, disk, and CPU usage of the current backup server.

1.1.11.3 View the backup status

The **Service Status** function 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.

Procedure

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

4. On the State page, view the current backup status.

- View the numbers of backup items that are in process, completed, timed out, and failed in the current system.
- View the statuses of the latest backup items of the current product.

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

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

1.1.12 NOC

1.1.12.1 Overview

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

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

1.1.12.2 Dashboard

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

Procedure

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

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

Item	Description	
	Ports Usage	<ul style="list-style-type: none"> · Ports Utilization: the proportion of ports in use to the total ports in the network devices. · Error Packets by Port (Top 5): the total number of error packets generated by device ports within a certain time range, of which the top 5 are displayed.
	Configuration Management	<ul style="list-style-type: none"> · Automatic Backup: the backup of startup configurations for all network devices. · Configuration Sync: the synchronization of running configurations and startup configurations for all network devices.
Network Monitoring	Alerts	The total number of alerts generated by network devices.
	Alerting Devices	The number of network devices that generate alerts and the total number of network devices.
	Alarm Details	The details of the alert.
Traffic Dashboard	SLB Overview	The bandwidth utilization of SLB clusters.

Item	Description	
	XGW Overview	The bandwidth utilization of XGW clusters.

1.1.12.3 Network topology

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

Procedure

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



Note:

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

- **Green:** The connection works properly.
- **Red:** The connection has an error.
- **Grey:** The connection is not enabled.

5. Click Detail in the upper-right corner to view the Device Properties and Port Status.

1.1.12.4 Resource management

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

1.1.12.4.1 Network elements

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

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

1.1.12.4.1.1 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.

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



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 and alert information of a device.**
 - a. **Click a device name, or click View in the Details column at the right of a device.**
 - b. **Under Port, view the port list, port working status, and other link information of the device.**
 - c. **Under Alert Info, view the alert information of the device.**

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

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

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

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

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

No search results exist if the device has a configuration exception or does not generate any logs during the selected time range.
5. Optional: You can filter the search results based on the log keyword.
6. Optional: Click Export to CSV in the upper-right corner to export the search results to your local computer.

1.1.12.4.1.1.3 Collection settings

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

1.1.12.4.1.1.3.1 Configure the collection interval

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

Procedure

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

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

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

1.1.12.4.1.1.3.2 Modify the collection interval

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

Procedure

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



Note:

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

5. Click Submit.

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

1.1.12.4.1.1.3.3 Add an OOB network segment

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

Context

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

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

Then, you can search for the IP address range of the managed device on this page.

Procedure

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

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

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

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

1.1.12.4.1.1.3.4 View the OOB network segment information

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

Procedure

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



Note:

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

1.1.12.4.1.2 Modify the device password

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

Procedure

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

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

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

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

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



Note:

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

1.1.12.4.1.3 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 Apsara Stack Operations.*

2. In the left-side navigation pane, choose NOC > Resource Management > Network Elements.

3. Click the Config Comparison tab.

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

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

5. Select the device and then click Compare Configuration.

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

1.1.12.4.2 Server Load Balancers

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

1.1.12.4.2.1 View the cluster monitoring information

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

Procedure

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

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.

1.1.12.4.2.2 View the instance monitoring information

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose NOC > Resource Management > Server Load Balancers.
3. Click the Instance Monitoring tab.
4. Select the cluster where the instance that you are about to view is located from the drop-down list. Enter the lb-id or VIP address that you are about to search for in the field and then click Search.
5. In the search result, view the monitoring information of the instance.
 - The first section is the basic information of the SLB instance, which allows operations engineers to troubleshoot problems and confirm the owner where a device belongs.
 - The second section is the operating water level graph of the instance. Select a time range and then click Search or select 5MIN, 30MIN, 1H, or 6H in the Quick Query section to view the operating water level graph of the instance in a specific time range, including the detailed bps and pps.

1.1.12.4.3 Collect IP addresses

The system regularly collects the IP addresses of all the physical networks in the current Apsara Stack environment based on the configured collection interval. You can search for the information of devices and ports to which a network segment or IP address belongs based on the network segment/IP address and subnet mask.

Procedure

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

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



Note:

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

1.1.12.4.4 IP address ranges

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

1.1.12.4.4.1 Import the planning file

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

Prerequisites

The IP address allocation list is obtained from [Apsara Stack Deployment Planner](#).

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose NOC > Resource Management > IP Address Ranges.
3. Click Import in the upper-right corner.
4. In the displayed dialog box, click Browse and then select the IP address allocation list.
5. Click Import.

1.1.12.4.4.2 Manually add the IP address pool information

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

Procedure

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

5. Click Add.

1.1.12.4.4.3 Modify the IP address pool information

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

Procedure

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



Note:

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

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

1.1.12.4.4.4 Export the IP address pool information

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

Procedure

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

1.1.12.4.4.5 Delete the IP address pool information

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

Procedure

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

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

1.1.12.5 Alert management

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

1.1.12.5.1 View and process current alerts

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **NOC > Alert Management > Alert Dashboard**.
3. Click the **Current Alerts** tab.
4. Enter a keyword in the search box in the upper-right corner and then click **Search**.
Alerts that meet the search condition are displayed.
5. Optional: You can filter the search results by device name, device IP address, or alert name.
6. Click **Details** in the Details column at the right of an alert to view the detailed alert information.
7. Find the reason why the alert is triggered and then process the alert.
 - If the alert does not affect the system normal operation, you can click **Ignore** in the Actions column to ignore the alert.
 - If the alert is meaningless, you can click **Delete** in the Actions column to delete the alert.

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

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

1.1.12.5.2 View history alerts

You can view history alerts on the History Alerts tab.

Procedure

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

1.1.12.5.3 Add a trap

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

Procedure

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

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

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

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

Configuration	Description	Example
Trap Msg	The message of the alert item.	<p>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.</p> <p>Configure the value strictly according to the device document. You cannot customize this value.</p>
Alert Type	Indicates whether this alert is of the fault type or the event type.	-
Association	<p>Indicates whether this alert has an event alert.</p> <p>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.</p>	-

5. Then, click Submit.

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

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

1.1.12.5.4 View a trap

You can view a trap configured in the current system.

Procedure

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



Note:

After the search results are displayed, you can click **Export to CSV** in the 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.

1.1.12.6 Network reconfiguration

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

1.1.12.6.1 Physical network integration

Physical Network Integration allows network operations engineers to perform automated integration of physical networks in Apsara Stack Operations (ASO) by entering the integration parameters. Network Operation Center (NOC)

automatically generates and issues the configurations to specific devices and then automatically performs the network integration test.

Procedure

1. *Log on to Apsara Stack Operations.*

2. In the left-side navigation pane, choose **NOC > Network Reconfiguration > Physical Network Integration.**

3. Enter the project name and then click **Create** to create a project.

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

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

5. Click **Next**.

6. Select a device.

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

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

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

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

c) Click **Save Project** in the upper-right corner to save the information of devices added to the **Target Device** section.

7. Click **Next**.

8. Configure the interface parameters.

a) In the **Configure Interfaces** step, click **Edit**.

b) Complete the parameter configurations and then click **Add** to add the interface to the list.

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

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

9. Click Next.

10. Configure the route parameters.

- a) **In the Configure Routes step, click Edit.**
- b) **Complete the parameter configurations and then click Add to add the route to the list.**
You can click Manage > Edit or Manage > Delete in the list to modify or delete the route.
- c) **Click Save Project in the upper-right corner to save the information.**

11. Click Next.

12. Configure the route policies.

- a) **In the Configure Route Policies step, click Edit.**
- b) **Complete the parameter configurations and then click Add to add the route policy to the list.**
You can click Manage > Edit or Manage > Delete in the list to modify or delete the route policy.
- c) **Click Save Project in the upper-right corner to save the information.**

13. Click Next.

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.

1.1.12.6.2 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*.

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose NOC > Network Reconfiguration > ASW Scale-up.
3. Select devices to be implemented.
 - a) In the Select Device step, enter a device name in the search box of the Devices on Live Network section and then click Search.
After adding a device, you can click Reset to clear the search condition and then search for and add another device.
 - b) Click Add at the right of the device to be implemented for this change to add the device on live network to the Target Device list.
To remove a device, click Manage > Delete in the Target Device list. You can also modify the logon username and password of the device by clicking Manage > Set the username and password.
4. Click Next.
5. Disable the DSW ports.
 - a) In the Disable DSW Port step, click Port Settings at the right of the device to be implemented.
 - b) Disable the corresponding port and then click Implement.
 - c) In the displayed dialog box, click OK to run the script commands.
6. Click Next.

7. Configure the DSW ports.

- a) **In the Configure DSW Port step, click Edit at the right of the device to be implemented. The Interface Parameter Configuration list is displayed.**
- b) **Select the Display Ports, enter the Port Description, IP Address, and Subnet Mask, and then click Add to add the interface parameter to the list.**
Then, you can click Manage > Edit or Manage > Delete to modify or delete the interface parameter.
- c) **After adding the interface parameter, click Implement at the right of the device.**
- d) **In the displayed dialog box, click OK to run the script commands.**
If an exception occurs after the implementation, you can click Back to roll back to the version before the implementation.

8. Click Next.**9. Configure the BGP.**

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

10. Click Next.**11. In the Upload ASW Configurations step, upload the new ASW configuration.****12. Click Next.****13. Enable the DSW ports.**

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

14. Click Next.

15. Perform the scale-up test.

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

1.1.12.7 Fault check

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

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

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

1.1.12.7.2 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 Apsara Stack Operations.*
2. **In the left-side navigation pane, choose NOC > Fault Check > Leased Line Discovery.**

3. Select a discovery source.

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

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

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

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

4. Click Next.**5. Configure the discovery parameters.**

- a) **In the Configure Parameters step, click Edit. The Configure Parameters list is displayed.**

- b) **Enter the Link Name, Destination IP Address, Source IP, Discovery Interval, Discoveries, and Discovery Timeout, and then click Add to add the information to the list.**

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

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

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.

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

Network operations engineers must create a project file for this inspection.

Parameters related to the project are saved in the file and you can click Manage > Import in the History section to import the project information if needed.

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

6. Select devices for inspection.

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

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

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

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

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



Note:

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

7. Click Next.

8. Select check items.

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

The added check items are displayed on the right.

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

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

9. Click Next.

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.

1.1.12.7.4 Configuration baseline audit

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

Procedure

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

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

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

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

1.1.13 Full Stack Monitor

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

1.1.13.1 SLA

SLA allows you to view the current state, history data, instance availability, and product availability of each cloud product. You can view the current and history fault state of products to obtain the SLA values and unavailable events of product instances within a certain time period.

1.1.13.1.1 View the current state of a cloud product

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

Procedure

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

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

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

1.1.13.1.2 View the history data of a cloud product

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

Procedure

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

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

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

1.1.13.1.3 View the availability of an instance

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

Procedure

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

1.1.13.1.4 View the availability of a product

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

Procedure

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

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

1.1.13.2 Operations full link logs

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

Context

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

Procedure

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



Note:

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

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

6. Select Abnormal logs only to only display the abnormal logs.

If code `! = 200`, `success=false`, or `error` exists in a log, the log is an abnormal log.

7. Enter a keyword in the search box to search for the related information in the search results.

8. Optional: After the search, you can click Export Log to export the search results to your local computer.

1.1.13.3 Correlation diagnosis and alarm

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

1.1.13.3.1 Full stack correlation alert

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

Procedure

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

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



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.

1.1.13.3.2 Server

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

Procedure

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

1.1.13.3.3 Network device

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm**.
3. Click the **Network Equipment** tab.
4. Enter the network device ID in the search box, select the time range, and then click **Search**. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
5. You can view the following information on this page.

- **Essential Information**

View the basic information of the network device.

- **Diagnostic Information**

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

- **Full Stack Alert**

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

1.1.13.3.4 ECS

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm.**
3. Click the ECS tab.
4. Enter the ECS instance ID in the search box, select the time range, and then click **Search.** You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
5. You can view the following information on this page.

- **Topology**

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

- **ECS Basic Info and HostNC Basic Info**

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

- **ECS Diagnosis Info and HostNC Diagnosis Info**

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

- **AVS Diagnosis and ECS-Alarm**

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

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

- **netdev**

View the traffic and packet information of the virtual NIC netdev on the host to which the ECS instance belongs. You can display the traffic or packet information by switching between the two tabs.

- **vport**

View the traffic, number of connections, and packet information of the virtual switch port vport on the host to which the ECS instance belongs. You can

display the traffic, number of connections, or packet information by switching among the tabs.

- **Network Alert Info**

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

- **Full Stack Alert**

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

1.1.13.3.5 RDS

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm.**
3. Click the 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.

- **Performance Data**

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

- **Diagnosis Info**

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

- **Network Alert Info**

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

- **Full Stack Alert**

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

1.1.13.3.6 SLB

You can use the 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.

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Full Stack Monitor > Correlation Diagnosis and Alarm.**
3. Click the SLB tab.

4. Enter the SLB instance ID in the search box, select the time range, and then click Search. You can select one hour, three hours, six hours, or one day as the time range, or customize the time range.
5. You can view the following information on this page: the topology of an SLB instance, the deployment information of an SLB cluster, the diagnosis information, the SLB bandwidth chart, the traffic of the LVS cluster, and the full stack correlation alert information.

- **Topology**

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

- **SLB Clusters**

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

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

- **Diagnostics**

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

- **SLB Bandwidth Chart**

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

- **Full Stack Alert**

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

1.1.13.3.7 VPC

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

Procedure

1. *Log on to Apsara Stack Operations.*

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

1.1.14 Pangu Operation Center

Pangu Operation Center displays the pangu grail, cluster information, node information, and pangu cluster status.

1.1.14.1 Pangu grail

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

Procedure

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

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

- **Overview**

Overview displays the storage space, server information, and health information of the selected product. Values of abnormal disks, abnormal

masters, abnormal chunk servers, and abnormal water levels in the Health section are displayed in red if they are larger than zero.

- **Heatmap of Health**

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

- **Data of Top 5 Services**

Data of Top 5 Services displays the data of the top 5 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.

1.1.14.2 Cluster information

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

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Pangu Operation Center > 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.

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

- **Overview** displays the storage space, server information, and health information of the selected cluster. Values of abnormal water levels, abnormal

masters, abnormal chunk servers, and abnormal disks in the Health section are displayed in red if they are larger than zero.

- **Alarm Monitor** displays the alert information of the selected cluster. You can perform a fuzzy search based on a keyword.
- **Replica** displays the replica information of the selected cluster.
- **Run Chart of Clusters** displays the charts of historical water levels, predicted water levels, number of files, number of chunk servers, and number of disks for the selected cluster.

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



Note:

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

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

1.1.14.3 Node information

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

Procedure

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

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

1.1.14.4 Pangu operation

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

Procedure

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

Clusters in different statuses are in different colors.

- Green indicates that the cluster works properly.
 - Yellow indicates that the cluster has a warning.
 - Red indicates that the cluster has an exception.
 - Dark red indicates that the cluster has a fatal error.
 - Grey indicates that the cluster is closed.
4. Move the pointer over a cluster name to view the service name, server name, and IP address to which the cluster belongs.

1.1.15 Task Management

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

1.1.15.1 Overview

Task Management has the following functions:

- Supports task overview and script overview to view task status and script status.
- Supports creating scripts and tasks.
- Supports the following four methods to run tasks: manual execution, timed execution, regular execution, and crontab.
- Supports the breakpoint function, which allows a task to stop between its two scripts and wait for manual intervention.
- Supports dynamically modifying execution parameters.
- Supports searching for tasks by name, status, and created time.

1.1.15.2 View the task overview

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

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose Task Management > Task Overview.

3. On the Task Overview page, you can perform the following operations:

- In the Dashboard section, view the number of tasks in the pending for intervention, running, failed, or completed state in the system.

Click the number of a certain state to go to the corresponding task page.

- In the Create Task section, create an operations task as follows:

a. Enter the task name and task description.

The task name cannot be the same as that of an existing task.

b. Select a target group.

A target group is the target of the task. You must select it by product > cluster > service > server role > virtual machine or physical machine.

c. Click .

d. Search for and add your uploaded script.

e. Click Start Now to run the 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 that require manual intervention in the Tasks To Be Intervened section.
- In the Running Tasks section, view tasks running in the last 24 hours.
- In the Create Script section, you can create a script by dragging a file to this section or clicking this section to upload a file.

You can upload a shell or Python script.

- In the Running Status in Last 7 Days section, view the running trend of tasks in the last seven days.

1.1.15.3 Create a script

Create an operations script and then you can run the script in a task.

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. In the left-side navigation pane, choose Task Management > Task Management.
3. Click the Scripts tab.
4. Click Create.
5. In the displayed dialog box, enter the script name and script description.

6. Click the Content tab on the right. Write scripts in the field.

Select a scripting language, bash or python, when you write scripts.

If you have a local script file, you can select a scripting language and then click Load Local File to load the information in the script to the field.

7. Click the Instructions tab on the right. Enter the script instructions in the field.

If you have the instructions included in a local file, click Load Local File to load the instructions to the field.

8. If you must depend on other files to run this script, click the Dependency File tab on the right. Click Select File to upload the dependency file.

9. Then, click Create.

You can view the created script in the script list.

1.1.15.4 Create a task

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

Prerequisites

A script is created.

Procedure

1. *Log on to Apsara Stack Operations.*

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

3. Click the Tasks tab.

4. Click Create.

5. In the displayed dialog box, enter the task name and task description, and then select a target group.

A target group is the target of the task. You must select it by product > cluster > service > server role > virtual machine or physical machine.

6. Add a script.

a) In the Scripts section, click .

b) In the displayed Create Script dialog box, search for and select the script to be created and then click  to add it to the script list on the right.

You can add multiple scripts according to the task requirement.

c) Click Create.

7. In the script list, configure the execution directory and execution parameter of the script, and then select whether intervention is required.

Configuration	Description
Execution Directory	The directory of the target (virtual machine or physical machine) that the script is running.
Execution Parameter	The parameters required when running the script . You can enter one or more parameters, with multiple ones separated by spaces.
Intervention Required	Select whether to enable the manual intervention. You are only required to select whether to enable the manual intervention if a task contains two or more scripts. If the manual intervention is enabled, the task stops and waits for manual intervention after you run the script.

8. Click Save.
9. In the displayed Execution Method dialog box, select the execution method of the task and then click OK.

You have four methods to run a task: manual execution, timed execution, regular execution, and crontab.

- **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.
- **Timed Execution:** Select the execution time. The task automatically runs when the time is reached.
- **Regular Execution:** Select the time interval and times to run the task. The task runs again if the execution condition is met.
- **crontab:** Configure the command to run the task periodically.

1.1.15.5 Modify a script

You can modify a script to update it.

Procedure

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

2. In the left-side navigation pane, choose Task Management > Task Management.
3. Click the Scripts tab.
4. Optional: Enter the script name and select the time range. Then, click Query.
5. Find the script to be modified and then click Modify in the Actions column.
6. Modify the content, instructions, and/or dependency file of the script, and then click Save.

1.1.15.6 Delete a script

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose Task Management > Task Management.
3. Click the Scripts tab.
4. Optional: Enter the script name and select the time range. Then, click Query.
5. Find the script to be deleted, and then click Delete in the Actions column.

1.1.15.7 View the execution status of a task

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose Task Management > Task Management.
3. Click the Tasks tab.
4. Optional: Enter the task name, select the task status, start date, and end date, and then click Query to search for tasks.
5. Find the task that you are about to view and then click Execution Status in the Actions column.
6. On the displayed Execution Status page, select a script and the target of the task to view the output.

1.1.15.8 Modify the execution parameters of a task

After a task is created, you can dynamically modify the execution parameters of the task.

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose Task Management > Task Management.
3. Click the Tasks tab.
4. Optional: Enter the task name, select the task status, start date, and end date, and then click Query to search for tasks.
5. Find the task to be modified and then click Modify in the Actions column.
6. On the Task Details tab, modify the execution parameters of each script in the task in sequence, and then click Save.

You can view the output and execution status of the task on the Export Script tab and Execution Records tab after the task runs.

1.1.15.9 Start a task

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose Task Management > Task Management.
3. Click the Tasks tab.
4. Optional: Enter the task name, select the task status, start date, and end date, and then click Query to search for tasks.
5. Find the task that you are about to start, and then click Start in the Actions column.

A message appears, indicating that the task is started. The task status changes from Not started to Running.

1.1.15.10 Delete a task

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

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose Task Management > Task Management.
3. Click the Tasks tab.
4. Optional: Enter the task name, select the task status, start date, and end date, and then click Query to search for tasks.

5. Find the task to be deleted and then click Delete in the Actions column.
6. In the displayed dialog box, click OK.

1.1.16 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 Apsara Stack Operations.*
2. In the left-side navigation pane, choose Task Management > Task Overview.
3. In the Tasks To Be Intervened section, find the task to be processed, and then click Confirm in the Actions column.

1.1.17 System Management

1.1.17.1 Overview

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

1.1.17.2 Department management

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

Context

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

Procedure

1. *Log on to Apsara Stack Operations.*

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

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

3. On this page, you can:

- **Add a department**

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

- **Modify a department**

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

- **Delete a department**



Notice:

Before deleting a department, make sure that no user exists in the department. Otherwise, the department cannot be deleted.

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

1.1.17.3 Role management

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

Context

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

Procedure

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

3. On the Role Management page, you can:

- Search for roles



Note:

To search for roles in ASO, you must have the ASO security officer role or system administrator role.

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

- Add a role



Note:

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

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

- Modify a role



Note:

To modify a role in ASO, you must have the ASO security officer role.

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

- Delete a role



Notice:

Before deleting a role, make sure that the role is not bound to any user. Otherwise, the role cannot be deleted.

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

1.1.17.4 Logon policy management

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

Context

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

Procedure

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

- Search for policies

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

- Add a policy

Click **Add Policy**. In the displayed dialog box, configure the **Policy Name**, **Start Time**, **End Time**, and 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.

1.1.17.5 User management

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

Prerequisites

Before you create a user, make sure that:

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

Context

User management provides different permissions for different users. During the system initialization, the system creates three default users: asosysadmin, asosecurity, and asoauditor. The default users are respectively bound to the following default roles: system administrator, security officer, and auditor officer. The permissions of these three roles are as follows:

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

Procedure

1. [Log on to Apsara Stack Operations.](#)
2. In the left-side navigation pane, choose System Management > Users. Click the Users tab.
3. On the Users tab, you can:
 - Search for users

**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/or 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.

1.1.17.6 Two factor authentication

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

Context

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

- Google two-factor authentication

This authentication method uses the password and mobile phone to provide double protection for accounts. You can obtain the logon key after 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 Apsara Stack Operations.*
2. In the left-side navigation pane, choose **System Management > Two Factor Authentication**.
3. On the **Two Factor Authentication** page, you can:
 - **Google two-factor authentication**
 - a. Select **Google Two-Factor Authentication as the Current Authentication Method**.
 - b. Click **Add User** in the upper-right corner. The added user is displayed in the user list.
 - c. Find the user that you are about to enable the Google two-factor authentication, and then click **Create Key** in the **Actions** column. After the key is created, you can click **Show Key** to display the key in plain text.
 - d. Enter the key in the Google authenticator app of your mobile phone. The app dynamically generates a verification code based on the time and key for logon. With the two-factor authentication enabled, you are required to enter the verification code on your app when logging on to the system.



Note:

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

- e. To disable the two-factor authentication, click **Delete Key** in the **Actions** column.
- **USB key authentication**
 - a. Select **USB Key Authentication as the Current Authentication Method**.
 - b. In the **Authentication Server Configuration** section, click **Add Server**. In the displayed dialog box, enter the IP Address and Port of the server, and then

click OK. The added server is displayed in the server list. Click Test to test the connectivity of the authentication server.

- c. In the User List section, click Add User. The added user is displayed in the user list.
- d. Find the user that you are about to enable the USB key authentication, and then click Bind Serial Number in the Actions column. In the displayed dialog box, enter the serial number to bind the user account with this serial number.



Note:

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

- e. Then, click Enable Authentication in the Actions column.
- **PKI authentication**
 - a. Select PKI Authentication as the Current Authentication Method.
 - b. In the Authentication Server Configuration section, click Add Server. In the displayed dialog box, enter the IP Address and Port of the server, and then click OK. The added server is displayed in the server list. Click Test to test the connectivity of the authentication server.
 - c. In the User List section, click Add User. Enter the Username, Full Name, and ID Card Number, and then click OK. The added user is displayed in the user list.
 - d. Find the user that you are about to enable the PKI authentication, and then click Bind in the Actions column. Enter the full name and ID card number of the user to bind the user account with the name and ID card number.
 - e. Then, click Enable Authentication in the Actions column.
 - **No authentication**

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

1.1.17.7 Application whitelist

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

Context

All the Apsara Stack Operations (ASO) services are accessed based on Operation Access Manager (OAM) permission management. Therefore, if your account does not have the corresponding role, your access requests are rejected. The 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 Apsara Stack Operations](#).
2. In the left-side navigation pane, choose System Management > Application Whitelist.
3. On the Application Whitelist page, you can:

- Add a whitelist

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

- Modify the permission

In the Permission drop-down list, modify the permission of the service to Read/Write or Read-only.

- Delete a whitelist

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

1.1.17.8 Server password management

Server Password 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, and/or IP address.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **System Management > Server Password**.

The Password Management tab displays the passwords of all the servers in the Apsara Stack environment.

3. On this tab, you can:

- **Search for servers**

On the Password Management tab, configure the product, hostname, and/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, and/or history IP address and then click Search to view the history of server password updates in the search results.

- **Show history passwords of servers**
 - a. **On the History Password tab, find a server.**
 - b. **Click Show in the Password column, and then the system displays the host password in plain text, which turns into cipher text after 10 seconds. Alternatively, directly click Hide to display the cipher text.**
- **View and modify the password configuration policy**

Click the Configuration tab. View the metadata, including the initial password, password length, and retry times, of server password management.

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

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

1.1.17.9 Operation logs

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

Context

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

Procedure

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

3. On the Log Management page, you can:

- Search for logs

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

- Delete logs

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

- Export logs

Click  to export the logs of the current page.

1.1.17.10 View the authorization information

The Authorization page allows customers, field engineers, or operations engineers to quickly view the service with an authorization problem and then troubleshoot the problem.

Prerequisites

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

If you are not an administrator-level user, a message indicating that you do not have sufficient permissions is displayed when you access this page.

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose System Management > Authorization.
3. View the authorization information on the Authorization Details page.



Note:

For formal authorization, you must enter the authorization code to view the authorization information. Obtain the authorization code in the authorization

letter attached by the project contract or contact the business manager (CBM) of your project to obtain the authorization code.

You can view the authorization information, including authorization version, customer information, authorization type, Elastic Compute Service (ECS) instance ID, the start date and end date of software license update and tech support, and service authorizations, of all services in the current Apsara Stack environment.

See the detailed authorization information and the corresponding description in the following table.

Authorization information	Description
Authorization Version	<p>You can use the BP number in the version to associate with a project or contract.</p> <ul style="list-style-type: none"> • TRIAL in the version indicates that the authorization is a trial one. The trial authorization is valid within 90 days from the date of deployment. • FORMAL in the version indicates that the authorization is a formal one. The authorization information of the service comes from the signed contract.
Authorization Type	Indicates the current authorization type and authorization status.
Customer information	Includes the customer name, customer ID, and customer user ID.
ECS Instance ID	The ECS instance ID in the Deployment Planner of the field environment.
Cloud Platform Version	The Apsara Stack version of the current cloud platform.
Authorization Created At	The start time of the authorization.

Authorization information	Description
Authorization information of a service	<p>Includes the service name, service content, authorization mode, service authorizations, software license update and tech support start date, software license update and tech support end date, and real-time authorization status.</p> <p>If the following information appears in the Authorization Status column of a service:</p> <ul style="list-style-type: none"> • RENEW Service Expired <p>Indicates that the customer must renew the subscription as soon as possible. Otherwise, the field operations services, including ticket processing, are to be terminated.</p> <ul style="list-style-type: none"> • Specifications Above Quota <p>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.</p>

1.2 Apsara Stack Doctor (ASD)

1.2.1 Apsara Stack Doctor introduction

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

Basic features

- **Provides data filtering, analysis, and processing for O&M data consumers.**
- **Provides encapsulation, orchestration, and rights management of O&M operations.**
- **Provides O&M experience accumulation and archiving capabilities.**

- **Provides troubleshooting, pre-diagnosis, health check, and early warning capabilities.**
- **Records O&M experience, prescriptions, monitoring data, and log data to support intelligent O&M.**

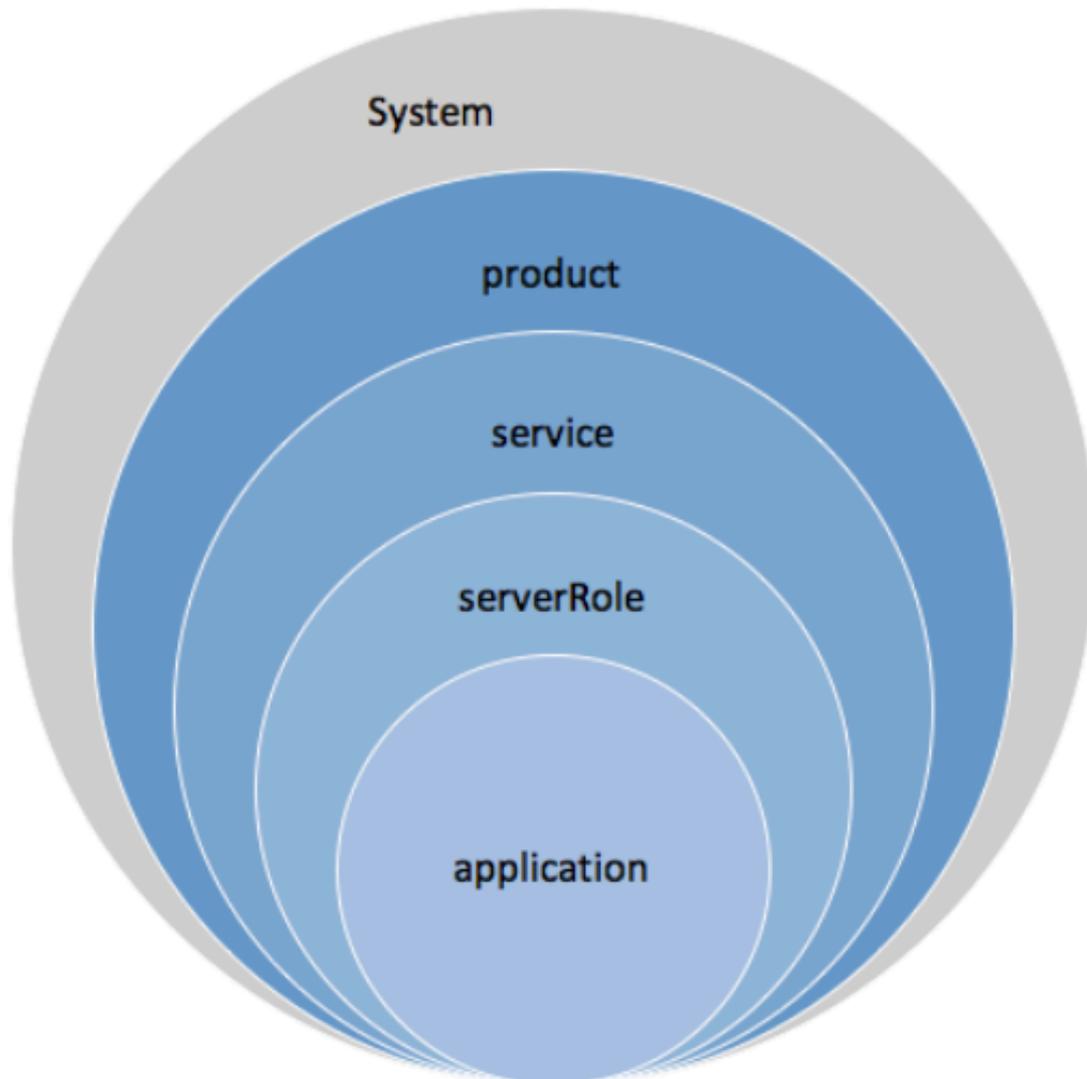
Benefits

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

Terms

Apsara Stack has five levels of release granularity, as shown in [Figure 1-2: Levels of release granularity](#).

Figure 1-2: Levels of release granularity



- **system**

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

- **product**

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

- **service**

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

- **server role (sr)**

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

- **application (app)**

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

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

1.2.2 Log on to Apsara Stack Doctor

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

Prerequisites

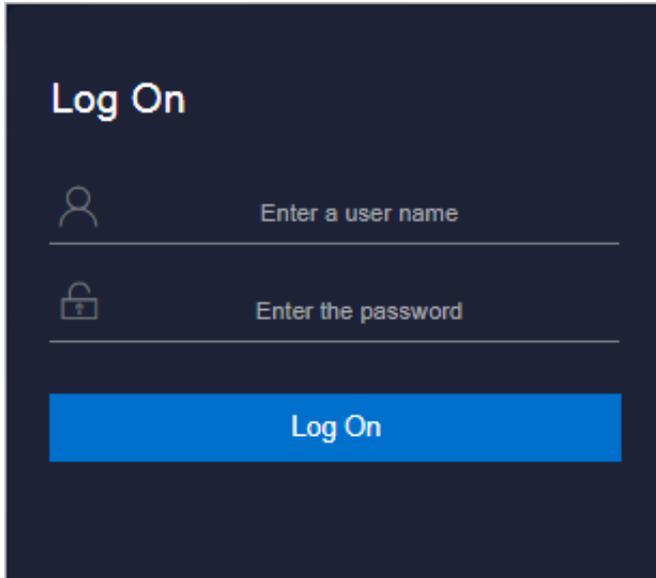
- ASO access address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

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.

Figure 1-3: Log on to ASO



Note:

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

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

6. In the Basic O&M region, click ASD.

1.2.3 Product dependency

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

Procedure

1. *Log on to Apsara Stack Doctor.*

2. To view the dependencies, perform the following operations:

- In the left-side navigation pane, choose **Dependencies > Product Dependency** to view the dependencies between different products.

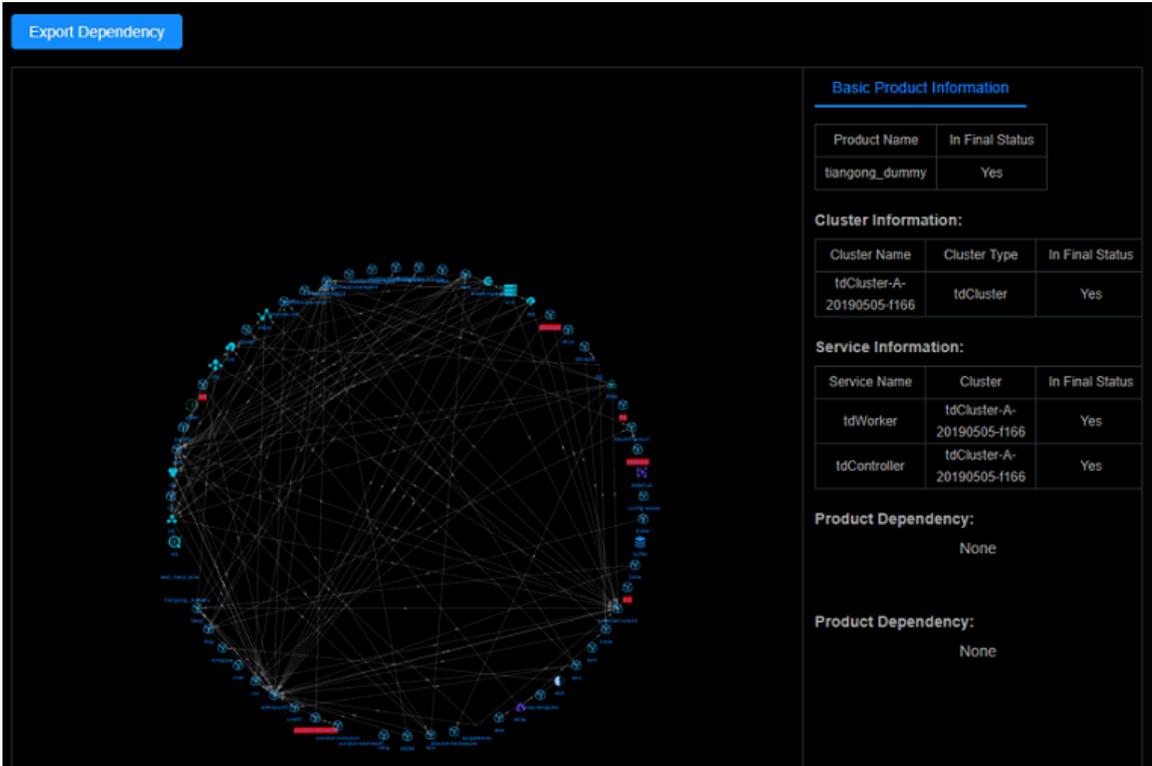
Click the name of a product, such as RDS, in the dependency graph.

Information about products that have dependencies with the selected product appears. Basic Product Information for the selected product appears on the right.



Note:

You can click **Export Dependency** in the upper-right corner of the page to export and save the product dependencies to your local machine.



Export Dependency

Basic Product Information	
Product Name	In Final Status
tiangong_dummy	Yes

Cluster Information:		
Cluster Name	Cluster Type	In Final Status
tdCluster-A-20190505-f166	tdCluster	Yes

Service Information:		
Service Name	Cluster	In Final Status
tdWorker	tdCluster-A-20190505-f166	Yes
tdController	tdCluster-A-20190505-f166	Yes

Product Dependency:
None

Product Dependency:
None

- In the left-side navigation pane, choose **Dependencies > Service Dependency** to view the dependencies between a service and its peripherals.

Select a cloud product and click a service in the dependency graph. You can view the information about other services that have dependencies with this service. Basic Service Information of this service appears on the right.

Cloud Product:

Basic Service Information

Product Name	base
Cluster Name	DataworksDqcCluster-A-20190423-7000
Cluster Type	DataworksDqcCluster
Service Name	bigdata-sre
Service In Final Status	Yes

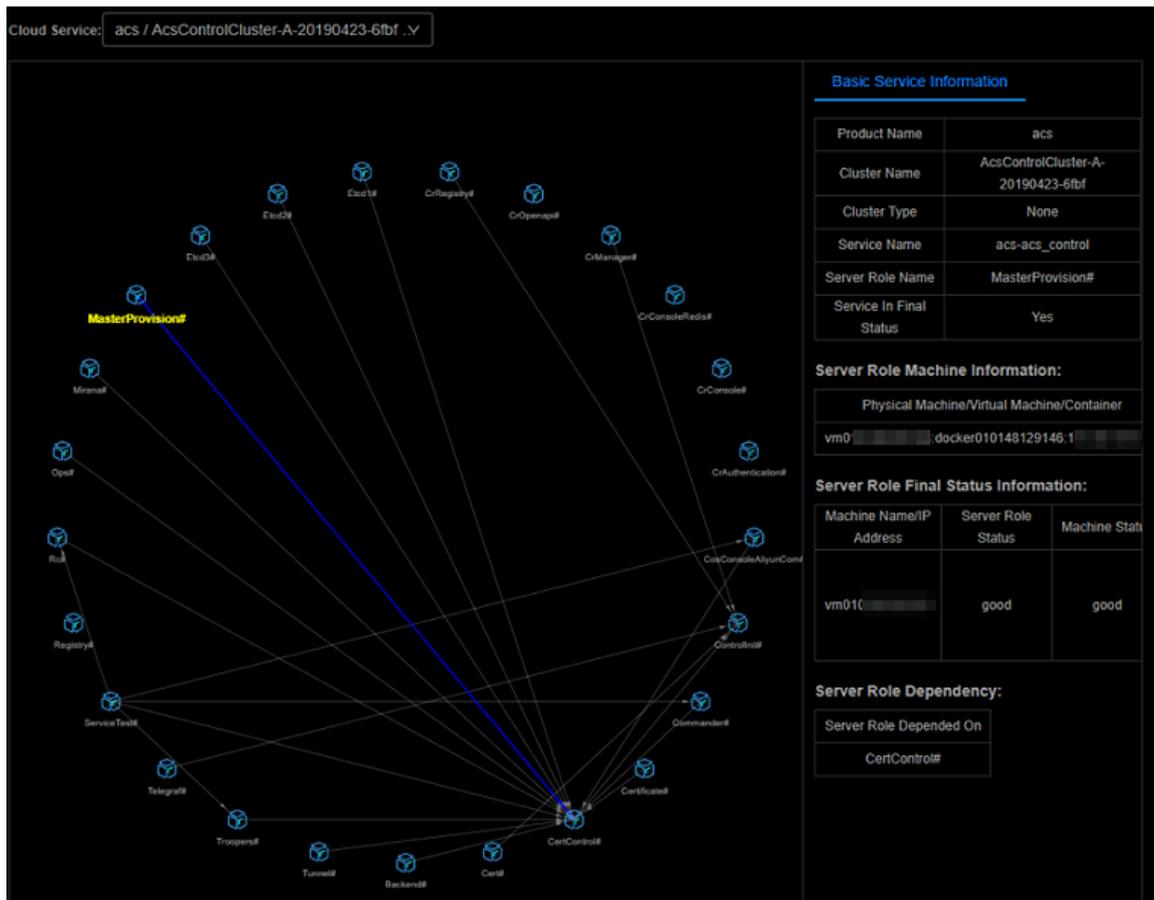
Server Role Information:

Server Role Name	In Final Status	Expected Machines In Final Status
Agent#	Yes	3

Service Dependency:
None

- In the left-side navigation pane, choose **Dependencies > Server Role Dependency** to view the dependencies between server roles.

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



1.2.4 Apsara Stack Inspection System

1.2.4.1 Apsara Stack Inspection System introduction

This topic gives a brief introduction about Apsara Stack Inspection System.

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

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

- **Big Data Inspection: inspects big data-related products.**

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



Notice:

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

1.2.4.2 Log on to Apsara Stack Inspection System

This topic describes how to log on to Apsara Stack Inspection System.

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.**

Figure 1-4: Log on to ASO

Log On

Enter a user name

Enter the password

Log On



Note:

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

3. Enter the correct username and password.

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

4. Click Log On to log on to ASO.

5. In the left-side navigation pane of the homepage, click Products. On the Product List page, click ASD.

6. In the left-side navigation pane, click Apsara Stack Inspection System.

1.2.4.3 Apsara Stack Inspection System overview

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

The Overview page consists of the following parts:

Quantity statistics

Statistics on the physical devices in Apsara Stack:

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



Physical machine statistics

The number of physical machines used for each Apsara Stack product.

Product	Physical Machines
acs	2
ads	15
base	6
blink	6
datahub	4
dataphin	3
dfs	6
drds	6
dts	4
ecs	89
ftp	1

Issues detected

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



Issues resolved

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



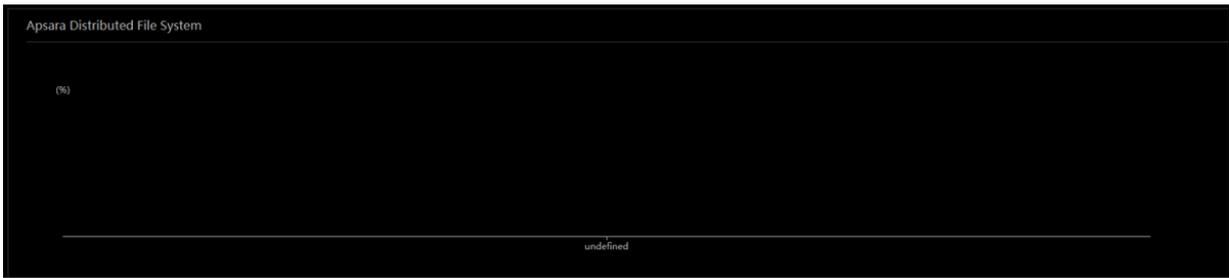
Utilization

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



Apsara Distributed File System

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



1.2.4.4 Platform Inspection

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

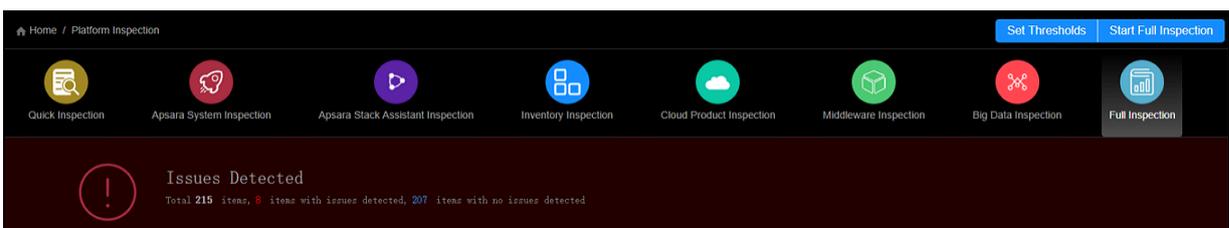
1.2.4.4.1 Overview

This topic gives a brief introduction about platform inspection.

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

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

When the inspection fails due to unknown reasons, you can click Reset Inspection in the upper-right corner to terminate this inspection.



1.2.4.4.2 Basic Inspection

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

Context

Basic Inspection consists of four parts:

- **check_ops**: checks the statuses of primary and secondary machines, including the health statuses of the memory modules, Docker containers, CPUs, and disks.

- **check_machine:** checks whether the machines are functioning normally.
- **check_pid:** checks the ratio of current process count to the maximum process count in the current system.
- **check_tianji:** checks whether the `service_manager`, `service_manager_supervisor`, `tj_master_main`, and `tj_proxy` processes exist on each machine in the cluster.
- **check_zombie process:** checks whether any zombie process exists on the machines in the Apsara Infrastructure Management Framework cluster.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. In the left-side navigation pane, choose Platform Inspection.
3. In the top navigation bar, click Basic Inspection.
4. Click Basic Inspection. The system starts the inspection and displays Inspecting...

If issues are detected during the inspection, Issues Detected appears after the inspection is completed. Items with issues detected and items with no issues detected appear at the lower part of the page.

5. If issues are detected, you can click the item with issues to view the details of the information.

1.2.4.4.3 Apsara System Inspection

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

Context

Apsara System Inspection consists of two parts:

- **Apsara Distributed File System inspection:** checks the Apsara Distributed File System metrics of each product, such as the read and write performance, data backup status, and version.
- **Apsara Name Service and Distributed Lock Synchronization System inspection:** checks the Apsara Name Service and Distributed Lock Synchronization System metrics of each product, such as the disk utilization and queuing status.



Notice:

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

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. **In the left-side navigation pane, choose Platform Inspection.**
3. **In the top navigation bar, click Apsara System Inspection.**
4. **Click Apsara System Inspection. The system starts the inspection and displays Inspecting...**

If issues are detected during the inspection, Issues Detected appears after the inspection is completed. Items with issues detected and items with no issues detected appear at the lower part of the page.

5. **If issues are detected, you can click the item with issues and move the pointer over detail in the prompt box to view the alert details.**

1.2.4.4.4 Inspection in Other Systems

ASA is called to perform Inspection in Other Systems.

Context

The Inspection in Other Systems module provides the following metrics:

- **Ip_conflict:** checks IP address conflicts in the current environment.
- **ntp:** This metric is used to check whether the system time of all machines (including Docker VMs and NCs) is synchronized with the NTP time. If not, the time offset (measured in ms) is reported.
- **rpm:** checks whether the RPM service of a machine is running properly.
- **dns_bind:** checks whether the IP address bound to a domain name is consistent with the requested IP address.
- **mem_quota:** checks the memory quota of a Docker host.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. **In the left-side navigation pane, choose Platform Inspection.**
3. **In the top navigation bar, click Inspection in Other Systems.**
4. **Click Inspection in Other Systems. The system starts the inspection and displays Inspecting...**

If issues are detected during the inspection, Issues Detected appears after the inspection is completed. Items with issues detected and items with no issues detected appear at the lower part of the page.

5. **If issues are detected, you can click the item with issues to view the alert details in the prompt box.**

1.2.4.4.5 Inventory Inspection

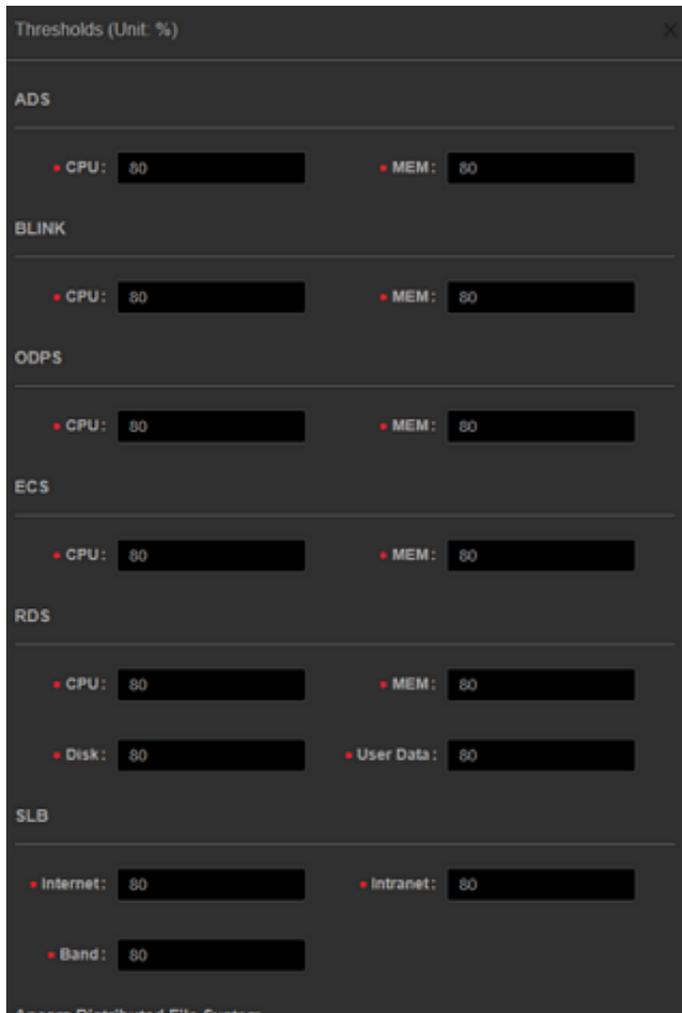
Inventory Inspection inspects the inventory status of ECS, OSS, RDS, SLB, , and each product's Apsara Distributed File System utilization, thereby calculating the remaining resources of each product.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. **In the left-side navigation pane, choose Platform Inspection.**

3. (Optional) Click Set Thresholds in the upper-right corner to set the CPU and memory thresholds of several products.

If the inventory of a product exceeds the threshold, an alarm is reported. The unit of thresholds is percentage (%). The current threshold is 80% by default. You can set the thresholds or use the default values as required.



4. In the top navigation bar, click Inventory Inspection.
5. Click Inventory Inspection. The system starts the inspection and displays Inspecting...

If issues are detected during the inspection, Issues Detected appears after the inspection is completed. Items with issues detected and items with no issues detected appear at the lower part of the page.

6. If issues are detected, you can click the item with issues and move the pointer over detail in the prompt box to view the alert details.

1.2.4.4.6 Cloud Product Inspection

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

Context

Cloud Product Inspection consists of three parts:

- **Component inspection:** checks for any abnormal server roles. For example, if a server role is in the upgrading state, it is not in the desired state.
- **SLB inspection:** checks whether the bandwidth of SLB exceeds the threshold and whether SLB has the coredump program.
- **RDS inspection:**
 - **Exception:** checks for any alerts reported for ApsaraDB Operations and Maintenance System.
 - **Task:** checks for interrupted tasks in ApsaraDB Operations and Maintenance System.

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. In the left-side navigation pane, choose Platform Inspection.
3. In the top navigation bar, click Cloud Product Inspection.
4. Click Cloud Product Inspection. The system starts the inspection and displays Inspecting...

If issues are detected during the inspection, Issues Detected appears after the inspection is completed. Items with issues detected and items with no issues detected appear at the bottom of the page.

5. If issues are detected, you can click the item with issues and move the pointer over `state_description` in the prompt box to view the alert details.

1.2.4.4.7 Middleware Inspection

The Middleware Inspection module inspects Butler-related items, such as middleware-dnscs, edas-customer, middleware, drds, dauthProduct, mq, csb, butler, and tlog.

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. In the left-side navigation pane, choose Platform Inspection.
3. In the top navigation bar, click Middleware Inspection.
4. Click Middleware Inspection. The system starts the inspection and displays Inspecting...

If issues are detected during the inspection, Issues Detected appears after the inspection is completed. Items with issues detected and items with no issues detected appear at the lower part of the page.

5. If issues are detected, you can click the item with issues and move the pointer over message in the prompt box to view the alert details.

1.2.4.4.8 Big Data Inspection

The Big Data Inspection module inspects big data services, including DataWorks, MaxCompute, PAI, StreamCompute, AnalyticDB, Dataphin, Elasticsearch, Graph Analytics, Quick BI, Apsara, DTBoost, AIMaster, and Apsara Big Data Manager.

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. In the left-side navigation pane, choose Platform Inspection.
3. In the top navigation bar, click Big Data Inspection.
4. Click Big Data Inspection. The system starts the inspection and displays Inspecting...

If issues are detected during the inspection, Issues Detected appears after the inspection is completed. Items with issues detected and items with no issues detected appear at the lower part of the page.

5. If issues are detected, you can click the item with issues to view the alert details in the prompt box.

1.2.4.4.9 Inspection Reports

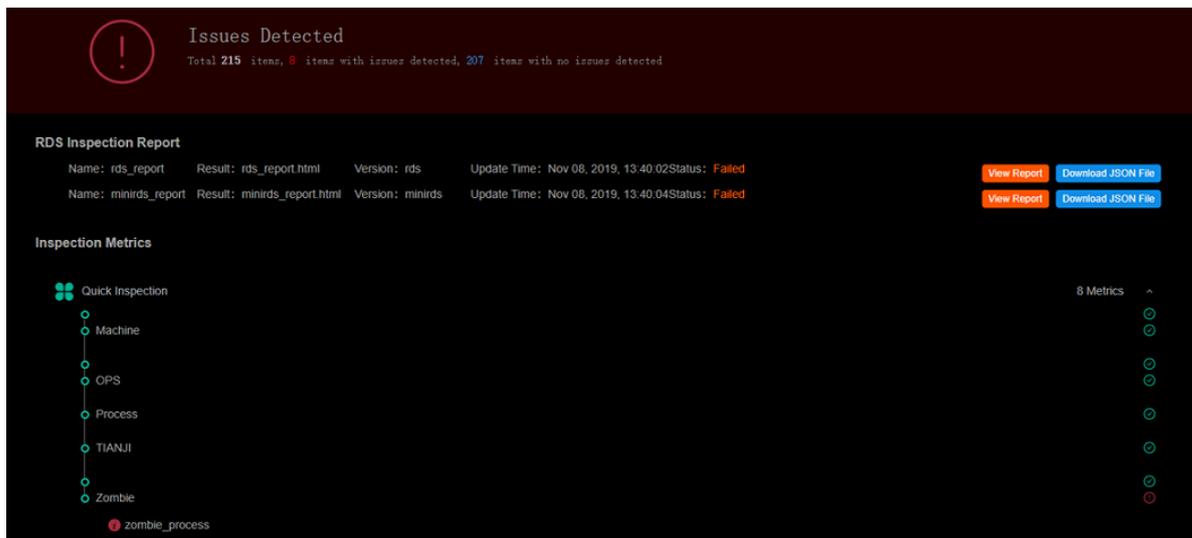
The Inspection Reports module summarizes inspection results by inspection type. You can view all detected issues on the Start All Inspections Reports page.

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. In the left-side navigation pane, choose Platform Inspection.

3. In the top navigation bar, click Start All Inspections Reports.

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



1.2.4.5 Inspection History

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

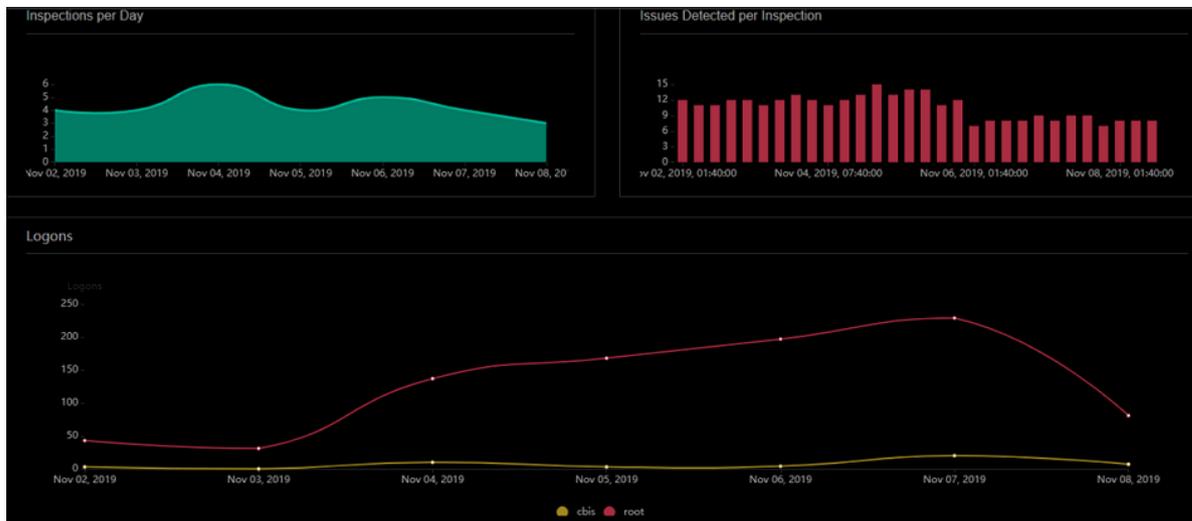
Procedure

1. *Log on to Apsara Stack Inspection System.*

2. In the left-side navigation pane, choose Inspection History.

You can view the following information:

- **Inspections per Day:** shows the number of inspections per day over a recent period of time.
- **Issues Detected per Inspection:** shows the number of issues detected per inspection.
- **Logons:** shows user logon statistics.
- **Detected Issues:** allows you to specify a time range and search the detected issues per inspection, and then view the detected issue details based on each inspection.



1.2.4.6 Work Reports

The Work Reports module records project information, issue tracking, abnormal utilization information, and daily reports.

1.2.4.6.1 Add project information

To facilitate routine maintenance and management, you can set Project Name, Version, Physical Devices, Products, and Technical Account Manager for a project.

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. In the left-side navigation pane, choose Work Reports.
3. In the Project Information section, click Edit.

4. Enter the information about the project.



Note:

When editing project information, you can click Cancel to cancel this edit.

5. Click Submit.

1.2.4.6.2 Issues

The Issues module displays issues that are encountered during routine maintenance.

1.2.4.6.2.1 Add issues

You can add issues to be tracked in routine maintenance to Work Reports, helping track the issue handling progress in a timely manner.

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. In the left-side navigation pane, choose Work Reports.
3. In the Issues section, click Add.
4. Enter the information of the issue to be tracked and click Submit.
5. **Optional:** To share the issue with others, you can select issues whose Status is Resolved and click Issue QR Code on the right of the Issues section. This way, you can send the issue to a DingTalk chatbot.

The DingTalk chatbot organizes the information and pushes it to the corresponding customers or O&M group.

1.2.4.6.2.2 Update the issue handling progress

After adding an issue, you can update the issue handling progress.

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. In the left-side navigation pane, choose Work Reports.
3. In the Issues section, find the row that contains the issue whose handling progress you want to update and click Edit in the Actions column.
4. In the Edit dialog box, modify the values of Status, Progress, and Resolution or Workaround.
5. Click Submit.

1.2.4.6.2.3 Remove issues

You can remove issues that you no longer need to follow.

Procedure

1. [Log on to Apsara Stack Inspection System](#).
2. In the left-side navigation pane, choose Work Reports.
3. In the Issues section, find the row that contains the target issue and click Remove in the Actions column.
4. In the dialog box that appears, click OK.

1.2.4.6.3 View resource utilization information

You can view the utilization information in the Resources, Apsara Distributed File System, Pangu Master, Physical Machines, and Docker Disk sections.

Procedure

1. [Log on to Apsara Stack Inspection System](#).
2. In the left-side navigation pane, choose Work Reports.
3. In the Current Utilization section, you can view the resource utilization information.

The following information is included:

- **Resources:** displays the total resource capacity and the percentage of used resources for each product.

When the absolute value of Growth Rate is greater than 5%, the system displays it in red.

- **Apsara Distributed File System:** displays the percentage of storage resources of Apsara Distributed File System used by each product.

When the storage utilization exceeds 75%, the system displays it in red.

- **Pangu Master:** displays the percentage of memory used by the pangu_master process of each product.

When the memory utilization exceeds 75%, the system displays it in red.

- **Physical Machines:** displays the disk utilization of a physical machine. You can specify the threshold.
- **Docker Disk:** displays the disk utilization of a Docker machine. You can specify the threshold.

4. Optional: To share the utilization information with others, click Threshold QR Code on the right of the Current Utilization section. This way, you can send the utilization information to a DingTalk chatbot.

The DingTalk chatbot organizes the information and pushes it to the corresponding customers or O&M group.

1.2.4.6.4 Reports

Reports contains three tabs: Issues Submitted Today, Unresolved Issues, and History. In the Reports section, you can view the issues submitted that day, unresolved issues, and issues resolved that month.

1.2.4.6.4.1 Add issues

You can add issues on the Issues Submitted Today tab based on the onsite O&M results to facilitate subsequent tracking and handling.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. **In the left-side navigation pane, choose Work Reports.**
3. **In the Reports section, click the Issues Submitted Today tab.**
4. **At the lower part of the issue list, click Add.**
5. **Set the parameters as required and click Save.**

1.2.4.6.4.2 Update issue progress

You can update issue progresses on the Issues Submitted Today, Unresolved Issues or History tab.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. **In the left-side navigation pane, choose Work Reports.**
3. **In the Reports section, click the Issues Submitted Today, Unresolved Issues, or History tab.**
4. **In the issue list, find the row that contains the issue whose progress you want to update and click Edit in the Actions column.**
5. **Update the values of Status, Handled By, and Resolved At.**
6. **After updating the information, click Save in the Actions column.**

1.2.4.6.4.3 Remove issues

You can remove issues that you do not need.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. In the left-side navigation pane, choose Work Reports.
3. In the Reports section, click the Issues Submitted Today, Unresolved Issues or History tab.
4. Find the issue that you no longer need and click Remove in the Actions column.
5. In the dialog box that appears, click OK.

1.2.4.6.4.4 View unresolved issues

On the Unresolved Issues tab, you can view the current unresolved issues.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. In the left-side navigation pane, choose Work Reports.
3. In the Reports section, click the Unresolved Issues tab.
4. In the issue list, view unresolved issues.

1.2.4.6.4.5 View reports

On the History tab, you can view the daily reports, weekly reports, and monthly reports.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. In the left-side navigation pane, choose Work Reports.
3. In the Reports section, click the History tab.
4. Switch among the Daily Report, Weekly Report, and Monthly Report tabs to view the corresponding data graphs.

1.2.4.6.5 Generate summary reports

After you complete all the inspection tasks at least once on the same day, you can generate a summary report of the issues that require attention and send the report out.

Prerequisites

Ensure that you complete Quick Inspection or Platform Inspection at least once on the same day.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. In the left-side navigation pane, choose Work Reports.
3. Click Generate Summary Report in the upper-right corner.
4. Set the issue statuses of Threshold Issues and Platform Issues to either Followed or Ignored.



Note:

Only the issues that are in the Followed state appear in the summary report.

5. Enter the summary of the report.
6. Click Summary Report QR Code & Import Issues to generate the QR code for the report.
7. This way, you can send the information to a DingTalk chatbot.

The DingTalk chatbot organizes the information and pushes it to the corresponding customers or O&M group.

1.2.4.6.6 Download work reports

After completing all inspection tasks at least once on the same day, you can download a daily report in EXCEL for onsite personnel.

Prerequisites

Ensure that you complete Quick Inspection or Platform Inspection at least once on the same day.

Procedure

1. *Log on to Apsara Stack Inspection System.*
2. In the left-side navigation pane, choose Work Reports.
3. Click Download EXCEL File in the upper-right corner to download the work report to your local machine.

The downloaded work report contains information, such as usage information, and a summary of important issues detected by the monitoring platforms of Apsara Stack.

1.2.4.7 Products

The Products module allows you to view the current resource usage of each product in Apsara Stack.

Context

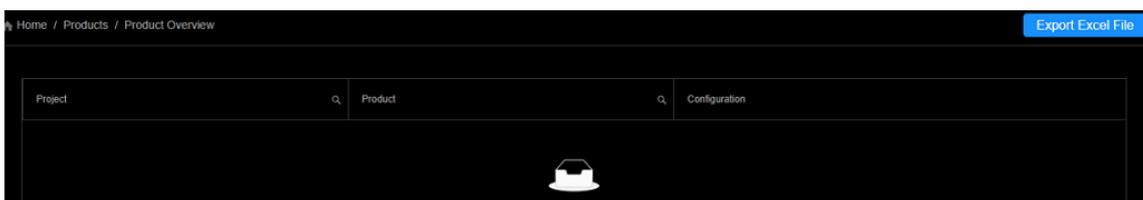
Information can be displayed using either of the following methods:

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

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. Perform the following operations to view the current resource usage:
 - In the left-side navigation pane, choose Products > Product Overview. On the Product Overview page, view the resource information for each product.

You can click Export Excel File to export the information as a spreadsheet.



- In the left-side navigation pane, choose Products > Product Category > Product Name to view the resource details of a product.

For example, choose Products > Database Services > RDS to view the resource details of ApsaraDB for RDS.

1.2.4.8 E2E ECS Links

E2E ECS Links allows you to search logs for issues in ECS instances.

Context

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

Procedure

1. [Log on to Apsara Stack Inspection System.](#)
2. In the left-side navigation pane, choose E2E ECS Links.

3. In the search bar, select the product name and service name. Then, enter the service instance ID or request ID, and click Run.

After you click Run, the system starts E2E ECS Links inspection. A line of information appears in the Tasks section.

In the Status column, different colors represent different statuses:

- Green: Inspected
 - Red: Failed to be inspected
 - Blue: Inspecting
4. Optional: After the status changes to Inspected, click Download. The log file in JSON format is downloaded to your local machine.



Note:

You can download logs only when the status is Inspected.

5. Optional: After the status changes to Inspected, click Learn More. On the Link Logs page, click the name of a log to view details.

Logs containing errors are marked with a red icon to make troubleshooting easier.

1.2.5 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.

1.2.5.1 RPM Check

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

Procedure

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

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

Host	Status
test_tianji_machine180	unavailable
ecsapigatewaylitetageb0a	unavailable
a36f04114.cloud.f05.amtest61	normal
vm010148064142	normal
vm010148064143	normal
vm010148064141	normal
vm010148064146	normal
a36f07206.cloud.f09.amtest61	normal
vm010148064026	normal
vm010148064023	normal

Table 1-2: Description of parameters on the RPM Check page

Parameter	Description
Host	The name of a host.
Status	<p>The status of a machine. Valid values:</p> <ul style="list-style-type: none"> normal: indicates that the machine is operating normally. unavailable: indicates that the machine is not operating normally or unavailable.

1.2.5.2 Virtual IP Check

The Virtual IP Check module allows you to obtain the virtual IP addresses that are incorrectly bound to IP addresses of backend services.

Procedure

1. *Log on to Apsara Stack Doctor.*

2. In the left-side navigation pane, choose ASA > Virtual IP Check.

Virtual IP Address	Virtual Port	Port	Backend IP Address	Cluster	Service	Server Role	Status
██████████	9090	21069	██████████	ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal
██████████	9090	21069	██████████	ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal
██████████	9090	21069	██████████	ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal
██████████	9090	21069	██████████	ots-ssd-A-20190423-6fc8	TableStore	TableStore.OTSFrontServer#	abnormal

Table 1-3: Parameters on the Virtual IP Check page

Parameter	Description
Virtual IP Address	The virtual IP address.
Virtual Port	The port corresponding to a virtual IP address.
Port	The port corresponding to the IP address of a backend service.
Backend IP Address	The IP address of a backend service.
Cluster	The cluster to which the IP address of a backend service belongs.
Service	The service to which the IP address of a backend service belongs.
Server Role	The server role to which the IP address of a backend service belongs.
Status	<p>The health status, indicating whether the binding between the virtual IP address and the IP address of the backend service is normal.</p> <ul style="list-style-type: none"> • normal: indicates that the virtual IP address is correctly bound to the IP address of the backend service. • abnormal: indicates that the virtual IP address is not bound to the backend IP address properly.

1.2.5.3 Volume Check

The Volume Check module allows you to view the volume details of Docker hosts.

Procedure

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

Container ID	Container Name	Host IP Address	Path	Disk Quota	Total Partition Space	Partition Space Used	Directory Space Used
2edb931eb098	bcc-api.Controller__controller.1558621857		/opt/backup_minirds	{"/": "40g"}	20G	1.1G	4.0K
2edb931eb098	bcc-api.Controller__controller.1558621857		/apsarapangu/disk8	{"/": "40g"}	45G	5.3G	4.0K
2edb931eb098	bcc-api.Controller__controller.1558621857		/apsarapangu	{"/": "40g"}	45G	5.3G	16K

Table 1-4: 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 <code>df</code> 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 <code>du</code> command.

1.2.5.4 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.

Procedure

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

Host	Time Offset
a36f04114.cloud.f05.amtest61	0
vm010148064142	0
vm010148064143	0
vm010148064141	0
vm010148064146	0

Table 1-5: Parameters on the NTP Check page

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

1.2.5.5 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	Type	Virtual Host
				

Table 1-6: 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.

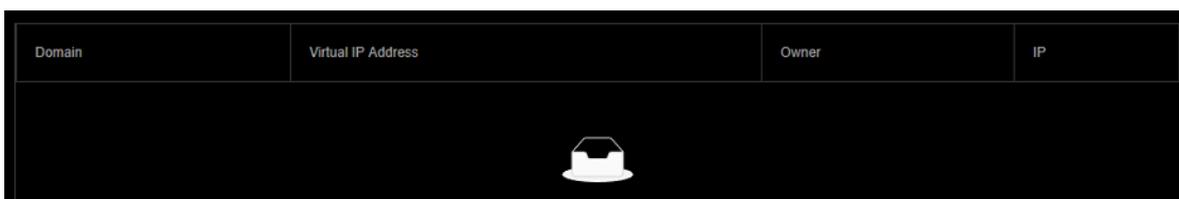
Parameter	Description
Server Role	The server role that requests the resource.
Type	The IP address type. Valid values: docker, vm, and physical.
Virtual Host	The hostname of the Docker virtual machine.

1.2.5.6 DNS Check

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

Procedure

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



Domain	Virtual IP Address	Owner	IP
			

Table 1-7: Parameters on the DNS Check page

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

1.2.5.7 IP Details

The IP Details module allows you to check the details of all IP addresses in the current environment, including the IP addresses of physical machines, Docker machines, and virtual machines, as well as virtual IP addresses.

Procedure

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

IP	Virtual Host	Type	Physical Host	Server Role
[REDACTED]		vip		Server Role Information
[REDACTED]		vip		Server Role Information
[REDACTED]		vip		Server Role Information
[REDACTED]		vip		Server Role Information
[REDACTED]		vip		Server Role Information

Table 1-8: Parameters on the IP Details page

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

3. Move the pointer over Server Role Information in the Server Role column to view server role details.

1.2.5.8 Quota Check

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

Procedure

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

Memory		CPU	Disk	
Container ID	Container Name	Container Memory	Hostname	Host Memory
cb88159341f2a	dtdream-dtcenter.Uim__uim.1559285092	4294967296	a36f04015.cloud.f04.amtest61	540732784640
c86de87d8d79c	vm010148065213	8643411968	a36f04015.cloud.f04.amtest61	540732784640
3eeee420a444c	asrbr-heimdallr.Heimdallr__heimdallr.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.

1.2.5.9 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.

1.2.5.10 Versions

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

Procedure

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

1.2.6 Support tools

1.2.6.1 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.

Physical Machine Name	Health Score	Host Address	Script Execution Status	Actions
a36f07203.cloud.f09.amtest61			Tunnel Error	View Report View Result Download Report
a36g03007.cloud.g03.amtest61			Not Executed	View Report View Result Download Report
a36f01060.cloud.f01.amtest61			Not Executed	View Report View Result Download Report
a36f01031.cloud.f01.amtest61			Not Executed	View Report View Result Download Report
a36f01002.cloud.f01.amtest61			Not Executed	View Report View Result Download Report
a36f01211.cloud.f03.amtest61			Not Executed	View Report View Result Download Report

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.

1.2.6.2 Use Support Tools

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

Procedure

1. *Log on to Apsara Stack Doctor.*

2. In the left-side navigation pane, choose Support Tools > Support Tools.

3. **Optional:** Select the target service, enter the host name or IP address, and click Search. The search results appear in the section below.

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	<p>Checks the following information:</p> <ul style="list-style-type: none"> • 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	<p>Collects and analyzes system logs, including the following operations:</p> <ul style="list-style-type: none"> • Collects information about the OS, network, disk, and hardware. • Diagnoses and analyze system logs. • Generates reports.

Diagnostic item	Description
oss_examine	Diagnoses OSS.

- 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	Search by hostname or IP address	Search	Run Diagnostics	Version: beta20190513	Upload File
HostName	ClusterName	IP Address	Diagnostics Execution Status	Executed At	Actions
a36f04013.cloud.f04.amtest61	ECS-I07River-A-6ffe	██████████	Succeeded	May 22, 2019, 15:49:49	Run Diagnostics View Report Download Report
a36f04011.cloud.f04.amtest61	ECS-I07River-A-6ffe	██████████	Succeeded	May 22, 2019, 15:49:49	Run Diagnostics View Report Download Report
a36f01109.cloud.f02.amtest61	ECS-I07River-A-6ffe	██████████	Succeeded	May 22, 2019, 15:49:49	Run Diagnostics View Report Download Report
a36f04210.cloud.f06.amtest61	ECS-I07River-A-6ffe	██████████	Succeeded	May 22, 2019, 15:49:48	Run Diagnostics View Report Download Report

- After the diagnostics are complete, click View Report in the Actions column to view the diagnostic result.
- Optional: After the diagnostics are complete, click Download Report in the Actions column to download the diagnostic results to your local machine.

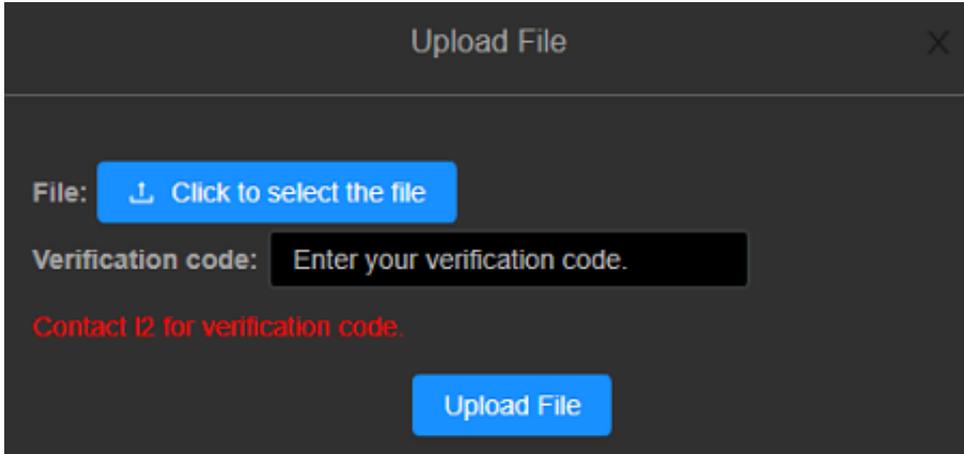
1.2.6.3 Update Support Tools

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

Procedure

- Log on to Apsara Stack Doctor.
- In the left-side navigation pane, choose Support Tools > Support Tools.
- 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.



1.2.6.4 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.

Procedure

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.

- 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.

Product:	Admin Gateway	IP	Diagnostics Execution Status	Executed At	Actions
ebs	vm010148000153		Not Executed	--	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.

1.2.6.5 Upload script files for EDAS diagnostics

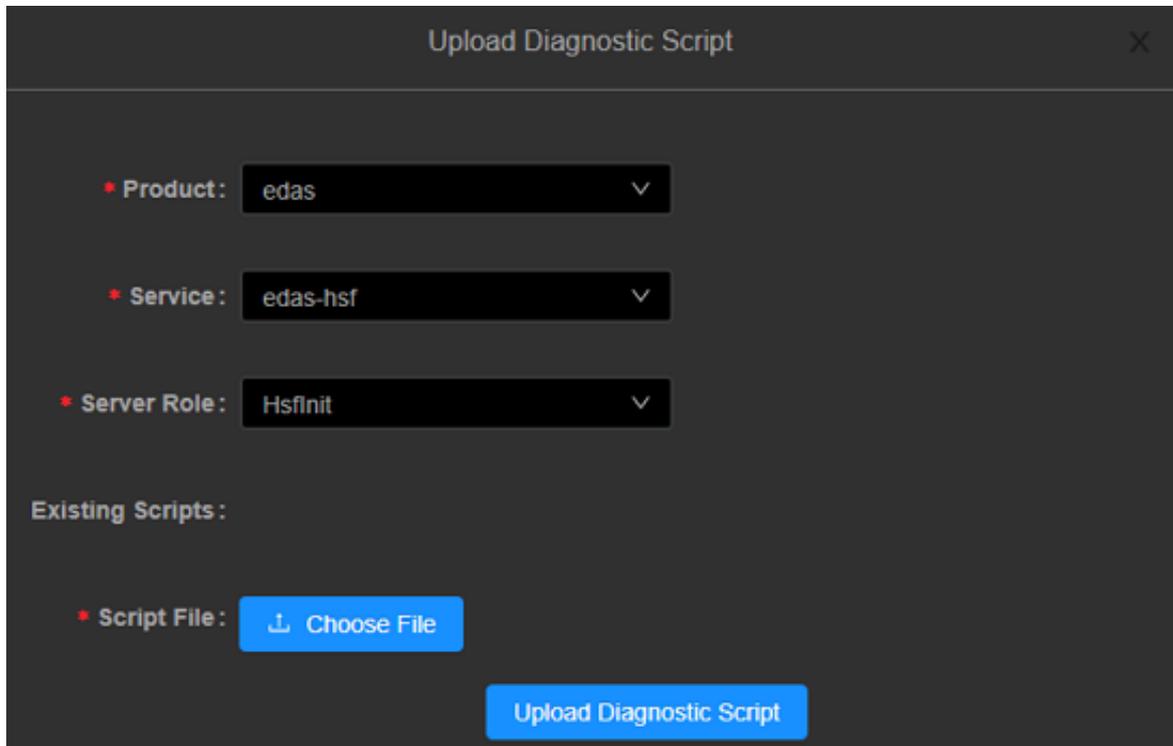
Before the diagnostics, you can unload script files to be executed for server roles.

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.



The screenshot shows a dark-themed dialog box titled "Upload Diagnostic Script". It features three dropdown menus for selection: "Product" (selected: edas), "Service" (selected: edas-hsf), and "Server Role" (selected: Hsfinit). Below these is a section labeled "Existing Scripts" which is currently empty. At the bottom, there is a "Script File" field with a "Choose File" button and an "Upload Diagnostic Script" button.

5. Click Choose File. In the dialog box that appears, select the script file to be uploaded. Click Open to add the script file to be uploaded.
6. Click Upload Diagnostic Script.

1.2.6.6 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.

Product:	edas	Service:	Select an item.	Refresh	Run Diagnostics	Upload Diagnostic Script
	Service	Server Role	Diagnostic Status	Cause of Failure	Actions	
<input type="checkbox"/>	edas-edasService	EdasServer	Diagnostics Failed	--	Run Diagnostics Download Report	
<input type="checkbox"/>	edas-edasService	CaIFs	Diagnostics Failed	No configuration snapshot.json	Run Diagnostics Download Report	
<input type="checkbox"/>	edas-edasService	EagleeyeConsole	Not Run	--	Run Diagnostics Download Report	
<input type="checkbox"/>	edas-edasService	EdasEam	Not Run	--	Run Diagnostics Download Report	

5. **After the tasks are completed, you can click Download Report in the Actions column corresponding to the server role to download the original diagnostic information.**

1.2.7 Service Availability

1.2.7.1 View Service Availability

Service Availability allows you to view the availability statuses of cloud services in Apsara Stack.

Context

It is used to verify the continuity of these cloud services.

During the hot upgrade of a service, you can use Service Availability to check whether the upgrade causes a service interruption, helping you detect and solve problems in a timely manner.

Procedure

1. *Log on to Apsara Stack Doctor.*
2. **In the left-side navigation pane, choose Service Availability > Service Availability.**

3. In the search bar, select the service you want to view and click Search to view its service status.

The following table describes the service statuses.

Service status	Description
Pending	The service availability inspection is not enabled for this service.
UNKNOWN	The service availability status of the service is unknown.
ERROR	The service availability status of the service is abnormal.
OK	The service availability status of the service is normal.



1.3 Operation Access Manager (OAM)

1.3.1 OAM introduction

Overview

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

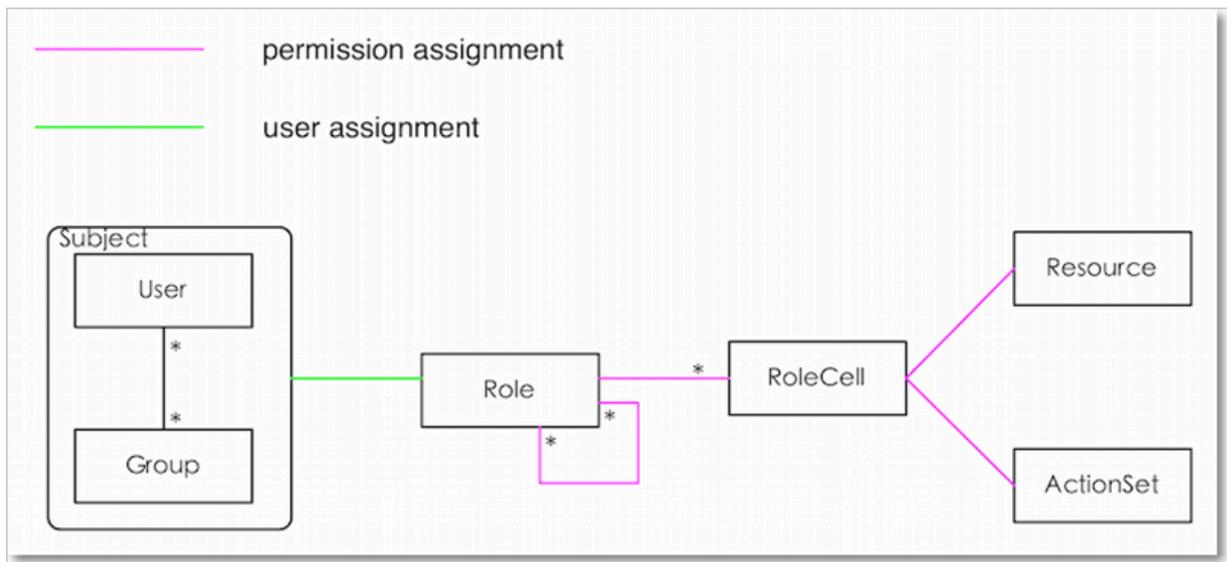
OAM permission model

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

, the frequency of role permission changes is less than that of user permission changes, simplifying the user permission management and reducing the system overhead.

See the *OAM permission model* as follows.

Figure 1-5: Permission model



1.3.2 Instructions

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

subject

Operators of the access control system. OAM has two types of subjects: users and groups.

user

Administrators and operators of operations systems.

group

A collection of users.

role

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

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

RoleHierarchy

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

RoleCell

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

resource

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

ActionSet

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

available authorizations

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

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



Note:

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

1.3.3 Quick start

This topic describes how to add and assign roles quickly.

1.3.3.1 Log on to OAM

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

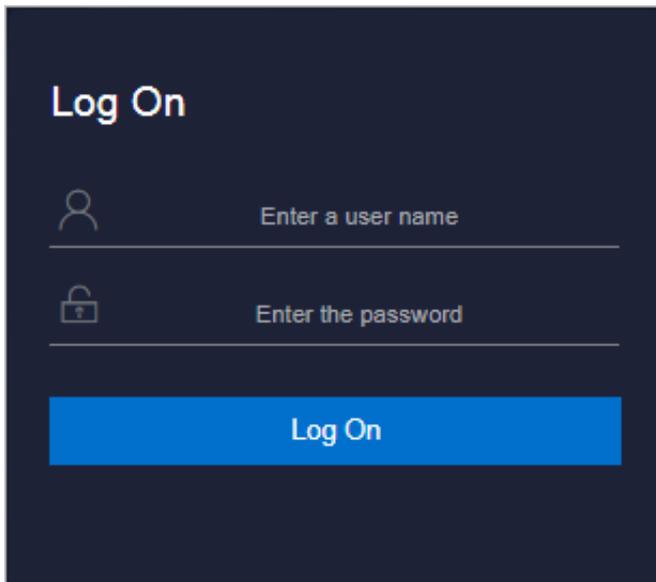
Prerequisites

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

Procedure

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

Figure 1-6: Log on to ASO



Note:

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

3. Enter the correct username and password.

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

4. Click Log On to log on to ASO.**5. In the left-side navigation pane, select Products.****6. Click OAM under Apsara Stack O&M.**

1.3.3.2 Create a group

Create a user group for centralized management.

Procedure

1. *Log on to OAM.*
2. **In the left-side navigation pane, choose Group Management > Owned Groups.**
3. **In the upper-right corner, click Create Group. In the displayed dialog box, enter the Group Name and Description.**
4. **Then, click Confirm.**

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

1.3.3.3 Add group members

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

Procedure

1. *Log on to OAM.*
2. **In the left-side navigation pane, choose Group Management > Owned Groups.**
3. **Find the group and then click Manage in the Actions column.**

4. Click **Add Member** in the **Group Member** section.
5. Select the search mode, enter the corresponding information, and then click **Details**. The user details are displayed.

Three search modes are available:

- **RAM User Account:** Search for the user in the format of *RAM username@primary account ID*.
- **Account Primary Key:** Search for the user by using the unique ID of the user's cloud account.
- **Logon Account Name:** Search for the user by using the logon name of the user's cloud account.

6. Click **Add**.
7. You can repeat the preceding steps to add more group members.

To remove a member from the group, click **Remove** in the **Actions** column at the right of the member.

1.3.3.4 Add group roles

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

Prerequisites

- The role to be added is created. For more information about how to create a role, see [Create a role](#).
- You are the owner of the group and the role.

Procedure

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

To remove a role from the group, click **Remove** in the **Actions** column at the right of the role in the **Role List** section.

1.3.3.5 Create a role

Procedure

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

The created tags are displayed in the dotted box.
 - e) Click Confirm to create the tags.
6. Click Confirm to create the role.

1.3.3.6 Add inherited roles to a role

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

Prerequisites

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

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

Procedure

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

1.3.3.7 Add resources to a role

You must add resources to a created role.

Procedure

1. [Log on to OAM](#).
2. In the left-side navigation pane, choose **Role Management > Owned Roles**.
3. Find the role and then click **Manage** in the **Actions** column.
4. Click the **Resource List** tab.
5. Click **Add Resource**.
6. Complete the configurations. For more information, see [Table 1-9: Configurations](#).

Table 1-9: Configurations

Configuration item	Description
BID	The deployment region ID.
Product	The cloud product to be added, for example, rds.  Note: The cloud product name must be lowercase. For example, enter rds, instead of RDS.
Resource Path	For more information about resources of cloud products and operations platforms, see 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**.

1.3.3.8 Add authorized users to a role

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

Prerequisites

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

Procedure

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

Four search modes are available:

- **RAM User Account:** search in the format of *RAM username@primary account ID*.
- **Account Primary Key:** search by using the unique ID of the user's cloud account.
- **Logon Account Name:** search by using the logon name of the user's cloud account.
- **Group Name:** search by group name.



Note:

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

7. Configure the expiration time.

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

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

To cancel the authorization, click Remove at the right of the authorized user on the Authorized Users tab.

1.3.4 Manage groups

Group Management allows you to view, modify, or delete groups.

1.3.4.1 Modify the group information

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

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, choose Group Management > Owned Groups.
3. Find the group and then click Manage in the Actions column.
4. Click Modify in the upper-right corner.
5. In the displayed Modify Group dialog box, modify the Group Name and Description.
6. Click Confirm.

1.3.4.2 View group role details

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

Prerequisites

A role is added to the group.

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, choose Group Management > Owned Groups.
3. Find the group and then click Manage in the Actions column.
4. In the Role List section, click Details at the right of a role.

5. On the Role Information page, you can:

- **Click the Inherited Role tab to view the information about the inherited roles.**
To view the detailed information of an inherited role, click **Details** in the **Actions** column at the right of the inherited role.
- **Click the Resource List tab to view the resource information of the role.**
To add other resources to this role, see [Add resources to a role](#).
- **Click the Inheritance Tree tab to view the basic information and resource information of the role and its inherited roles by using the inheritance tree on the left.**

1.3.4.3 Delete a group

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

Prerequisites

The group to be deleted does not contain members.

Procedure

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

1.3.4.4 View authorized groups

You can view the groups to which you are added on the **Authorized Groups** page.

Context

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

Procedure

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

1.3.5 Manage roles

Role Management allows you to view, modify, transfer, or delete roles.

1.3.5.1 Search for roles

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

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, choose Role Management > Owned Roles.
3. **Optional:** Enter the role name.
4. Click Search to search for roles that meet the search condition.



Note:

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

1.3.5.2 Modify the role information

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

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, choose Role Management > Owned Roles.
3. Find the role and then click Manage in the Actions column.
4. Click Modify in the upper-right corner.
5. In the displayed Modify Role dialog box, modify the Role Name, Description, Role Type, and Tag.
6. Then, click Confirm.

1.3.5.3 View the role inheritance tree

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

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, choose Role Management > Owned Roles.
3. Find the role and then click Manage in the Actions column.

4. Click the Inheritance Tree tab.

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

1.3.5.4 Transfer roles

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

Procedure

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

Four search modes are available:

- **RAM User Account:** search in the format of *RAM username@primary account ID*.
 - **Account Primary Key:** search by using the unique ID of the user's cloud account.
 - **Logon Account Name:** search by using the logon name of the user's cloud account.
 - **Group Name:** search by group name.
6. Click Transfer to transfer the roles to the user or group.

1.3.5.5 Delete a role

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

Prerequisites

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

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, choose Role Management > Owned Roles.
3. At the right of the role to be deleted and then click Delete.

1.3.5.6 View authorized roles

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

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, choose **Role Management > Authorized Roles**.
3. On the **Authorized Roles** page, you can view the name, owner, description, modified time, and expiration time of the role assigned to you.

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

1.3.5.7 View all roles

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

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, choose **Role Management > All Roles**.
3. On the **All Roles** page, view all the roles in the system.
You can search for roles by **Role Name** on this page.
4. At the right of a role, click **Details** to view the inherited roles, resources, and inheritance tree information of the role.

1.3.6 Search for resources

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

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, select **Search Resource**.
3. Enter the **Resource** and **Action** in the search boxes, and then click **Search** to search for roles that meet the conditions.
4. At the right of a role, click **Details** in the **Actions** column to view the inherited roles, resources, and inheritance tree information of the role.

1.3.7 View the personal information

You can view the personal information of the current user and test the permissions on the Personal Information page.

Procedure

1. *Log on to OAM.*
2. In the left-side navigation pane, select Personal Information.
3. In the Basic Information section, you can view the username, user type, created time, AccessKey ID, and AccessKey Secret of the current user.



Note:

Click Show or Hide to show or hide the AccessKey Secret.

4. In the Test Permission section, test if the current user has a certain permission.
 - a) Enter the resource information in the Resource field.



Note:

Use the English input method when entering values in the Resource and Action fields.

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

1.3.8 Appendix

1.3.8.1 Default roles and their functions

This topic describes the default roles in Operation Access Manager (OAM) and their functions.

1.3.8.1.1 Default role of OAM

This topic describes the default role of Operation Access Manager (OAM) and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations
Super administrator	An administrator with root permissions	*:*	*	10

1.3.8.1.2 Default roles of Apsara Infrastructure Management Framework

This topic describes the default roles of Apsara Infrastructure Management Framework and the corresponding available authorizations.

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

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

Role name	Role description	Resource	Actions	Available authorizations
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

1.3.8.1.3 Default role of DataQ - Smart Tag Service

This topic describes the default role of DataQ - Smart Tag Service and the corresponding available authorizations.

Role name	Role description	Resource	Actions	Available authorizations
Operations administrator of DataQ - Smart Tag Service	Performs all operations on the operations side of DataQ - Smart Tag Service	*:dtboost-ops:*	["operate"]	0

1.3.8.1.4 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
Webapp-rule read-only	Has the read-only permission to Webapp-rule projects, which allows you to view all the configurations and statuses	26842:webapp-rule:*	["read"]	0

1.3.8.1.5 Default roles of the workflow console

This topic describes the default roles of the workflow console and the corresponding available authorizations.

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

1.3.8.1.6 Default role of Tianjimon

This topic describes the default role of Tianjimon and the corresponding available authorizations.

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

1.3.8.2 Permission lists of operations platforms

This topic describes the permissions of operations platforms.

1.3.8.2.1 Permission list of Apsara Infrastructure Management Framework

This topic describes the permissions of Apsara Infrastructure Management Framework.

Resource	Action	Description
*:tianji:services:[sname]:tjmontemplates:[tplname]	delete	DeleteServiceTjmonTpl
*:tianji:services:[sname]:tjmontemplates:[tplname]	write	PutServiceTjmonTpl
*:tianji:services:[sname]:templates:[tplname]	write	PutServiceConfTpl
*:tianji:services:[sname]:templates:[tplname]	delete	DeleteServiceConfTpl
*:tianji:services:[sname]:serviceinstances:[sname]:tjmontemplate	read	GetServiceInstanceTjmonTpl
*:tianji:services:[sname]:serviceinstances:[sname]:tssessions	terminal	CreateTsSessionByService
*:tianji:services:[sname]:serviceinstances:[sname]:template	write	SetServiceInstanceTpl
*:tianji:services:[sname]:serviceinstances:[sname]:template	delete	DeleteServiceInstanceTpl
*:tianji:services:[sname]:serviceinstances:[sname]:template	read	GetServiceInstanceTpl
*:tianji:services:[sname]:serviceinstances:[sname]:tags:[tag]	delete	DeleteServiceInstanceProductTagInService

Resource	Action	Description
*:tianji:services:[sname]:serviceinstances:[siname]:tags:[tag]	write	AddServiceInstanceProductTagInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:resources	read	GetServerroleResourceInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	write	OperateSRMachineInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	read	GetMachineSRInfoInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines:[machine]	delete	DeleteSRMachineActionInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	read	GetMachinesSRInfoInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	delete	DeleteSRMachinesActionInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:machines	write	OperateSRMachinesInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:resources	read	GetAppResourceInService

Resource	Action	Description
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles:[serverrole]:apps:[app]:machines:[machine]:tianjilogs	read	TianjiLogsInService
*:tianji:services:[sname]:serviceinstances:[siname]:serverroles	read	GetServiceInstanceServerrolesInService
*:tianji:services:[sname]:serviceinstances:[siname]:schema	write	SetServiceInstanceSchema
*:tianji:services:[sname]:serviceinstances:[siname]:schema	delete	DeleteServiceInstanceSchema
*:tianji:services:[sname]:serviceinstances:[siname]:rollings:[version]	write	OperateRollingJobInService
*:tianji:services:[sname]:serviceinstances:[siname]:rollings	read	ListRollingJobInService
*:tianji:services:[sname]:serviceinstances:[siname]:resources	read	GetInstanceResourceInService
*:tianji:services:[sname]:serviceinstances:[siname]:machines:[machine]	read	GetMachineAllSRInfoInService
*:tianji:services:[sname]:serviceinstances:[siname]	write	DeployServiceInstanceInService
*:tianji:services:[sname]:serviceinstances:[siname]	read	GetServiceInstanceConf
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:name	read	GetMachineAppFileListInService
*:tianji:services:[sname]:serverroles:[serverrole]:machines:[machine]:apps:[app]:files:download	read	GetMachineAppFileDownloadInService

Resource	Action	Description
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps: [app]:files:content	read	GetMachineAppFileContentInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps: :[app]:filelist	read	GetMachineFileListInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps: [app]:dockerlogs	read	DockerLogsInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps: [app]:debuglog	read	GetMachineDebugLogInService
*:tianji:services:[sname]: serverroles:[serverrole]: machines:[machine]:apps	read	GetMachineAppListInService
*:tianji:services:[sname]: serverroles:[serverrole]: apps:[app]:dockerinspect	read	DockerInspect
*:tianji:services:[sname]: schemas:[schemaname]	write	PutServiceSchema
*:tianji:services:[sname]: schemas:[schemaname]	delete	DeleteServiceSchema
*:tianji:services:[sname]: resources	read	GetResourceInService
*:tianji:services:[sname]	delete	DeleteService
*:tianji:services:[sname]	write	CreateService
*:tianji:projects:[pname]: machinebuckets:[bname]: machines:[machine]	read	GetMachineBucketMachineInfo
*:tianji:projects:[pname]: machinebuckets:[bname]: machines	read	GetMachineBucketMachines

Resource	Action	Description
*:tianji:projects:[pname]: machinebuckets:[bname]	write	CreateMachineBucket
*:tianji:projects:[pname]: machinebuckets:[bname]	write	OperateMachineBucket Machines
*:tianji:projects:[pname]: machinebuckets:[bname]	delete	DeleteMachineBucket
*:tianji:projects:[pname]: machinebuckets:[bname]	read	GetMachineBucketMach inesLegacy
*:tianji:projects:[pname]: machinebuckets	read	GetMachineBucketList
*:tianji:projects:[pname]: projects:[pname]:clusters :[cname]:tssessions:[tssessionname]:tsses	terminal	UpdateTsSessionTssBy Cluster
*:tianji:projects:[pname]: projects:[pname]:clusters: [cname]:tssessions	terminal	CreateTsSessionByCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:tjmontemplate	read	GetServiceInstanceTj monTplInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:template	delete	DeleteServiceInstanc eTplInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:template	write	SetServiceInstanceTm plInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:template	read	GetServiceInstanceTm plInCluster
*:tianji:projects:[pname]:clusters:[cname]: serviceinstances:[siname]:tags:[tag]	write	AddServiceInstancePr oductTagInCluster

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

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

Resource	Action	Description
*: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	GetServiceInstanceServerrolesInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	delete	DeleteServiceInstanceSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:schema	write	SetServiceInstanceSchemaInCluster
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]:resources	read	GetInstanceResourceInCluster

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	delete	DeleteServiceInstance
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	write	CreateServiceInstance
*:tianji:projects:[pname]:clusters:[cname]:serviceinstances:[siname]	read	GetServiceInstanceConfInCluster
*:tianji:projects:[pname]:clusters:[cname]:rollings:[version]	write	OperateRollingJob
*:tianji:projects:[pname]:clusters:[cname]:rollings	read	ListRollingJob
*:tianji:projects:[pname]:clusters:[cname]:resources	read	GetResourceInCluster
*:tianji:projects:[pname]:clusters:[cname]:quota	write	SetClusterQuotas
*:tianji:projects:[pname]:clusters:[cname]:machinesinfo	read	GetClusterMachineInfo
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	read	GetMachineAllSRInfo
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	write	SetMachineAction
*:tianji:projects:[pname]:clusters:[cname]:machines:[machine]	delete	DeleteMachineAction
*:tianji:projects:[pname]:clusters:[cname]:machines	write	OperateClusterMachines
*:tianji:projects:[pname]:clusters:[cname]:difflist	read	GetVersionDiffList

Resource	Action	Description
*:tianji:projects:[pname]:clusters:[cname]:diff	read	GetVersionDiff
*:tianji:projects:[pname]:clusters:[cname]:deploylogs:[version]	read	GetDeployLogInCluster
*:tianji:projects:[pname]:clusters:[cname]:deploylogs	read	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
*:tianji:ids:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	CreateRackunit
*:tianji:ids:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	write	SetRackunitAttr
*:tianji:ids:[idc]:rooms:[room]:racks:[rack]:rackunits:[rackunit]	delete	DeleteRackunit
*:tianji:ids:[idc]:rooms:[room]:racks:[rack]	write	SetRackAttr
*:tianji:ids:[idc]:rooms:[room]:racks:[rack]	write	CreateRack

Resource	Action	Description
*:tianji:idcs:[idc]:rooms:[room]:racks:[rack]	delete	DeleteRack
*:tianji:idcs:[idc]:rooms:[room]	write	CreateRoom
*:tianji:idcs:[idc]:rooms:[room]	delete	DeleteRoom
*:tianji:idcs:[idc]:rooms:[room]	write	SetRoomAttr
*:tianji:idcs:[idc]	delete	DeleteIdc
*:tianji:idcs:[idc]	write	SetIdcAttr
*:tianji:idcs:[idc]	write	CreateIdc

1.3.8.2.2 Permission list of DataQ - Smart Tag Service

This topic describes the permission of DataQ - Smart Tag Service.

Resource	Action	Description
:dtboost-ops:	operate	Performs all operations on the operations side of DataQ - Smart Tag Service

1.3.8.2.3 Permission list of Webapp-rule

This topic describes the permissions of Webapp-rule.

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

1.3.8.2.4 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

Resource	Action	Description
26842:grandcanal	write	Restarts, retries, rolls back, and terminates a workflow

1.3.8.2.5 Permission list of Tianjimon

This topic describes the permission of Tianjimon.

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

1.4 Apsara Opsapi Management

1.4.1 Apsara Opsapi Management system overview

The Apsara Opsapi Management (Opsapi platform) system is a platform to manage operations and maintenance APIs and SDKs in the Apsara Stack environment in a unified manner. This system also manages API and SDK versions.

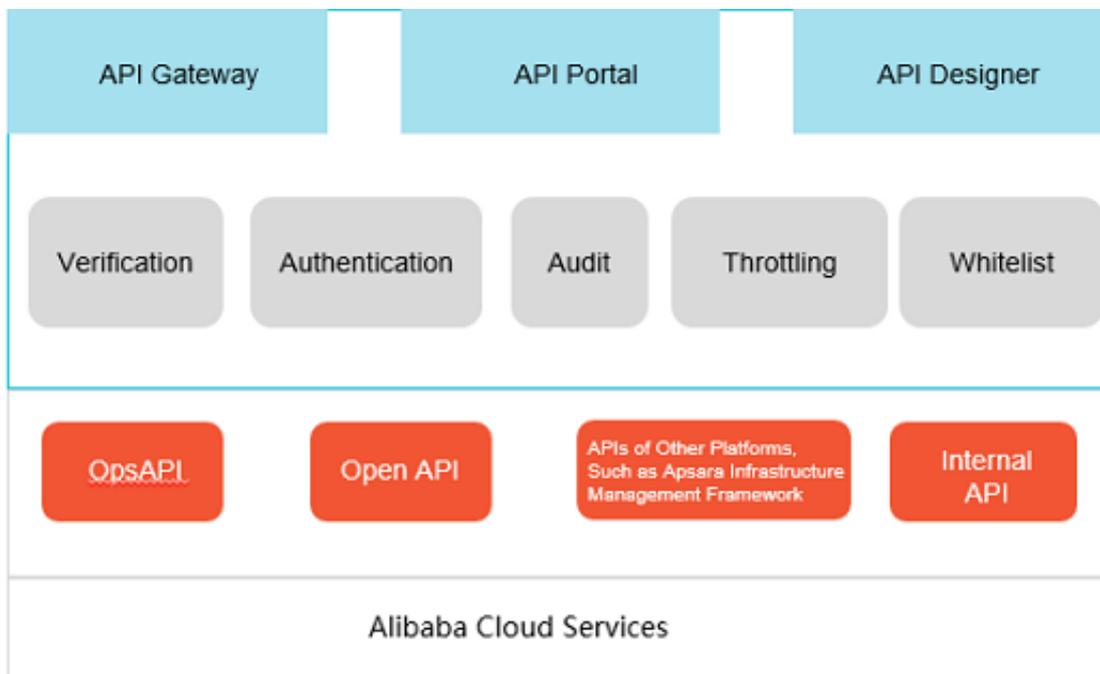
Most Apsara Stack products provide OpenAPIs that are used at the tenant side. Only a few Apsara Stack products provide APIs that are used at the operations and maintenance side. To address the business need at the operations and maintenance side and meet custom development requirements of users such as developing their own operations and maintenance console or obtaining operations and maintenance data, Alibaba Cloud provides the Apsara Opsapi Management system.

Features of the Apsara Opsapi 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 an API flow based on the existing API, which facilitates custom business.
- Manages versions and relationships between these versions. These versions include Apsara Stack versions, product versions, SDK versions, and API versions.
- Supports SDKs. The Apsara Opsapi Management system provides Java and Python SDKs to call operations and maintenance APIs.

Figure 1-7: Basic architecture of the Apsara Opsapi Management system shows the basic architecture of the Apsara Opsapi Management system.

Figure 1-7: Basic architecture of the Apsara Opsapi Management system



The Apsara Opsapi 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 Opsapis.
- **api-node:** the API designer.

1.4.2 Log on to the Apsara Opsapi Management system

The Apsara Opsapi Management system provides basic platform management functions for operations and maintenance engineers. These functions include API management, version management, test management, and system management. This topic describes how to log on to the Apsara Opsapi Management system.

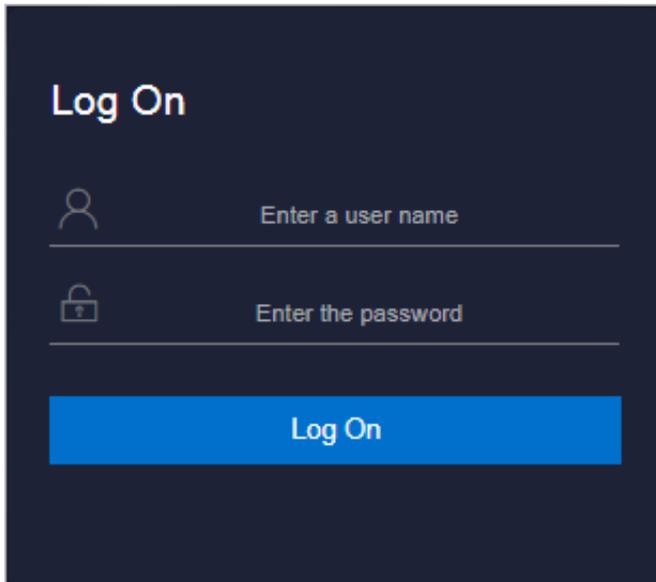
Prerequisites

- **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.

Figure 1-8: Log on to ASO

**Note:**

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

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

5. In the left-side navigation pane of the homepage, click **Products**. On the **Product List** page, click **Apsara Opsapi Management**.

1.4.3 API management

The Apsara Opsapi Management system provides APIs of various products in the Apsara Stack environment. You can manage these APIs, such as uploading, querying, editing, testing, and deleting APIs. You can also use the API designer to customize APIs.

1.4.3.1 Register APIs

An API can be defined in an XML file. Each API corresponds to one XML file. You can upload an XML file to register an API in the Apsara Opsapi Management system.

Context

The following fields must be defined in an XML file:

- API name
- Namespace (or product name)
- API type
- Parameter information

For more information about descriptions of XML files, see the POP document.

When you upload APIs, you can upload one or more XML files simultaneously. You need to compress the XML files as a ZIP file before you upload these files.

Procedure

1. [Log on to the Apsara Opsapi Management system](#).
2. In the left-side navigation pane, choose **API Platform > APIs**. The **APIs** page appears.
3. In the upper-right corner, click **Upload API**. Select the XML or ZIP file to be uploaded.

After you upload the XML or ZIP file, you can view the uploaded APIs in the API list on the **APIs** page.

1.4.3.2 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 Opsapi Management system.*
2. In the left-side navigation pane, choose **API Platform > APIs**. The APIs page appears.
3. **Optional:** Select a product from the **Select Product** drop-down list. Enter an API name in the search box.

Enter a full or partial API name to search for APIs.

4. Click the  icon in the Actions column corresponding to the API information about which is to be modified.
5. In the Edit API dialog box that appears, click the **Basic Information** tab, modify basic information, and click **Save**.

The following table describes parameters of an API.

Parameter	Description
API Name	The name of the API.
Type	The type of the API. Different products have different API types. The types are as follows: <ul style="list-style-type: none"> • opsAPI: the API for operations and maintenance • OpenAPI: the OpenAPI 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 the API.

6. Click the **Edit Parameters** tab to modify configuration information such as the request parameters, response parameters, and error handling mechanism. Follow these steps: Set the request parameters in the **Parameters** section, response parameters in the **ResultMapping** section, and error handling mechanism in the **ErrorMapping** section.
7. After the modification is completed, click **Save** to submit the modification results.

1.4.3.3 Test APIs

The Apsara Opsapi Management system allows you to test APIs online to check whether an API is available. During the test, you can save input parameters as a test case for subsequent execution.

Procedure

1. [Log on to the Apsara Opsapi Management system.](#)
2. In the left-side navigation pane, choose API Platform > APIs. The APIs page appears.
3. On the APIs page, click the  icon in the Actions column corresponding to the API to be tested.
4. In the Test API dialog box, set Request Parameters.

Request parameters may vary with APIs. The following table describes typical request parameters.

Parameter	Required	Description
regionId	Yes	The region ID of the test environment.
accessKeyId and accessKeySecret	Yes	The identification of the visitor. You can obtain it from the Apsara Stack console.

5. After request parameters are configured, click Send.

The Apsara Opsapi Management system sends a corresponding test request to the configured domain name. The response appears in the Responses section.

6. Optional: After the test is completed, click Save As Test Case for subsequent execution of this test case on the Test Cases page (choose Testing Platform > Test Cases).

1.4.3.4 Remove information about APIs

You can remove information about an API that you no longer need.

Procedure

1. [Log on to the Apsara Opsapi Management system.](#)
2. In the left-side navigation pane, choose API Platform > APIs. The APIs page appears.

3. Optional: Select a product from the Select Product drop-down list. Enter an API name in the search box.

Enter a full or partial API name to search for APIs.

4. Click the  icon in the Actions column corresponding to the API information about which is to be removed.

5. In the message that appears, click OK.

1.4.3.5 API design

The Apsara Opsapi Management system provides an API designer to help you customize APIs.

1.4.3.5.1 Designers

When Apsara Stack Opsapis do not match the APIs you are using or you need to customize APIs to meet requirements of specific projects, you can use an API designer to assemble and create desired APIs in the flow design process.

The API designer is built based on the open-source project Node-RED. Node-RED is a powerful tool launched by IBM to build Internet of Things (IoT) applications. It uses the visual programming method that allows developers to connect predefined code blocks (nodes) to perform tasks. Connected nodes are a combination of input nodes, processing nodes, and output nodes. When they are connected to form a flow, they are able to process requests such as HTTP requests.

Node-RED is highly capable of customizing flows and processing HTTP messages, which can be easily expanded.

To design an API, follow these steps:

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose API Platform > API Design. The API Designer page appears.
3. Drag and drop the components on the left to the flow chart section. You can combine these components to complete a specific flow chart and design an API.

1.4.3.5.2 Designer nodes

This topic describes several nodes that are used in the designer.

To customize Opsapis, the Apsara Opsapi Management system adds some nodes through the mechanism provided by Node-RED.

Typical nodes are described as follows:

- **api-request:** used to create a request to access an Opsapi. There is a small icon before each node. After you click the icon, a request is automatically 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:** detailed operations involved in a request process. For example, set protocol types of the request 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 returned data formats.

1.4.3.5.3 Design an API flow

Each customized API has its own API flow. Each API flow consists of connected nodes which include 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 some 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.

1.4.4 Version management

1.4.4.1 Apsara Stack version management

Apsara Stack has multiple versions that vary with projects.

1.4.4.1.1 Add information about versions

You can add information about Apsara Stack versions as needed to manage the relationships among Apsara Stack versions, products, and product versions.

Context

Each Apsara Stack version can have either one release version or one snapshot of the on-premises environment or deployment environment. It can be distinguished by its version name and description.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Versions > Apsara Stack Versions**. The **Apsara Stack Versions** page appears.
3. In the upper-right corner of the page, click **Add Version**.
4. In the **Add Version** dialog box that appears, set **Apsara Stack Version**, **Version**, and **Release Notes**.

We recommend that you enter information that is related to the current version for **Release Notes**.

5. Click **Submit**.

1.4.4.1.2 Select products for an Apsara Stack version

After adding information about an Apsara Stack version, you can select products that are supported in an Apsara Stack version based on version output conditions.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Versions > Apsara Stack Versions**. The **Apsara Stack Versions** page appears.

3. In the version list, click **Configure Products in the Product column** corresponding to the specified Apsara Stack version.
4. In the **Configure Products** dialog box that appears, select a version from the **Version drop-down list** and select the check box in the **Output column** corresponding to the product.

Configure Apsara Stack Products: 3.3
✕

Product	Version	Output
Ecs	3.1.0 ▼	<input checked="" type="checkbox"/>
Rds	3.1.0 ▼	<input checked="" type="checkbox"/>
Oss	2.4.2 ▼	<input type="checkbox"/>
Vpc	3.4.0 ▼	<input checked="" type="checkbox"/>
Slb	3.4.0 ▼	<input checked="" type="checkbox"/>

submit

5. Click **Submit** to generate information of the products of the specified Apsara Stack version.

1.4.4.1.3 Compare versions

You can use the version comparison function to compare the product differences between two Apsara Stack versions. Based on these product differences, you can further learn about the differences of their APIs as well as of the definitions and parameters of these APIs.

Context

- **Apsara Stack version:** Each Apsara Stack version can have either one release version or one snapshot of the on-premises environment or deployment environment. Versions are distinguished by its version name and description.
- **Product version:** the specific version of a product when each Apsara Stack version is released, such as RDS 3.7.0. An Apsara Stack release version can have only one version of a specific product.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Versions > Apsara Stack Versions**. The Apsara Stack Versions page appears.
3. In the upper-left corner of the page, click **Compare Versions**.
4. On the **Select Version** tab, select two versions to be compared, and click **Next**.

Select the source version	Select the 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
v3.9	v3.9

5. On the **Version Differences** tab, you can compare the product differences between the two versions. For example, you can view versions for which product information has been added or removed.

6. Click a product. Click Next to go to the Product Differences tab. You can compare the differences in product APIs between these two versions.

You can view functions for which APIs have been added or removed, and APIs remain the same in these two versions.

Compare Apsara Stack Versions X

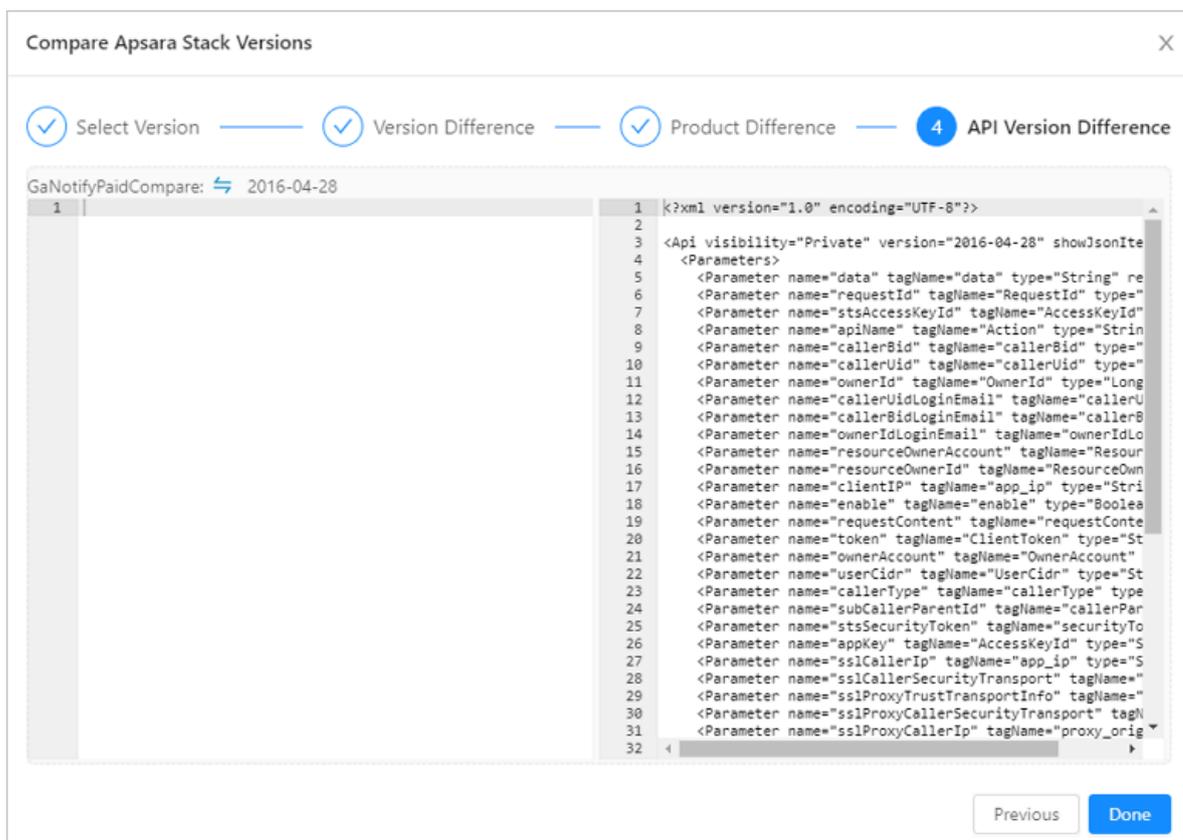
✓ Select Version
✓ Version Difference
3 Product Difference
 4 API Version Difference

ProductPaiSource Version: N/A, Target Version: 3.3.1
 New API: 801 Items, Deleted API: 0 Items, Changed API: 0 Items, Unchanged API: 0 Items

Function Name	Source Version	Target Version	Status	Remark
CountCloudInstances			2016-04-28	new ● New
DescribeNetworkQuotas			2016-04-28	new ● New
GaFillParams			2016-04-28	new ● New
GaFillProduct			2016-04-28	new ● New
GaNotifyPaid			2016-04-28	new ● New
GaOrderCheck			2016-04-28	new ● New
InnerAddBillingTag			2016-04-28	new ● New

Previous
Next

7. Click an API. Click Next to view the changes that are made to this API.



1.4.4.1.4 Remove information about Apsara Stack versions

If you no longer need the Apsara Stack version, you can remove its version and output information.

Procedure

1. [Log on to the Apsara Opsapi Management system.](#)
2. In the left-side navigation pane, choose **Versions > Apsara Stack Versions**. The **Apsara Stack Versions** page appears.
3. Click the  icon in the **Actions** column corresponding to the Apsara Stack version information about which is to be removed.
4. In the message that appears, click **OK**.

1.4.4.2 Product baseline management

Product baselines are a set of configurations used by Apsara Stack products to define products, services, service roles, and applications. The Apsara Opsapi Management system provides basic information about products, services, and service roles. During initialization, the Apsara Opsapi Management system automatically scans all product baseline information in the Apsara Stack

environment. You can use the system to scan the metadatabases and servers of services and service roles.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Versions > Product Baselines**. The **Product Baselines** page appears.
3. Follow these steps:
 - In the upper-right corner of the page, click **Scan Apsara Stack Environment** to scan metadatabases and servers that correspond to all products and update their information in the system.
 - Select a product from the drop-down list to query the service and service roles of the product.
 - Click the  icon in the Actions column corresponding to the service role to scan the metadatabases of the service role.
 - Select a service role. click the  icon in the Actions column corresponding to the service role to scan the server of the service role.

1.4.4.3 Product management

Operations and maintenance engineers can manage information of current Apsara Stack versions and product versions in real time.

Context

- **Apsara Stack version and product version:** Each Apsara Stack version can have only one specific version of products.
- **Product version and SDK version:** Each product version can have one SDK version
-

1.4.4.3.1 Add information about products

You can add information about a product you need to manage.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. In the upper-right corner of the page, click **Add Product**.

4. In the Add Product dialog box that appears, set Product Name and Product Description.
5. Click Submit to add information about a product.

1.4.4.3.2 Add information about product versions

After you add information about a product, you need to add its product version and API version information for subsequent version management.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
3. On the Products page, click the  icon in the Actions column corresponding to the product about which the version information is to be added.
4. In the Add Version dialog box that appears, set Version and API Version.

Parameter	Description
Version	The version of the current product.
API Version	The API version of the current product.

5. Click Submit.

1.4.4.3.3 Import information about APIs

You can import information about a preset API to the Apsara Opsapi Management system.

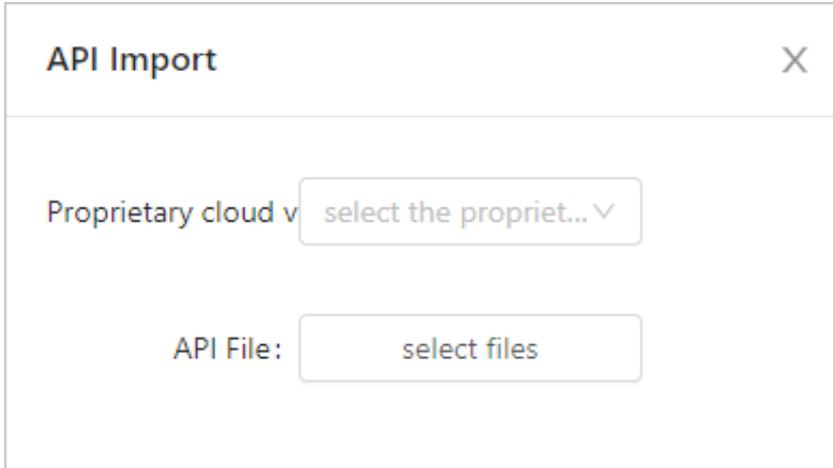
Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
3. In the product list, click the  icon in the Actions column corresponding to the product to be managed.
4. In the dialog box that appears, select the product version from the left drop-down list.

If information about the API has been imported for the product version, this API is displayed in the APIs section.

5. In the upper-right corner, click Import API.

- In the Import API dialog box that appears, select the Apsara Stack Version to which the API to be imported belongs and the corresponding API File.



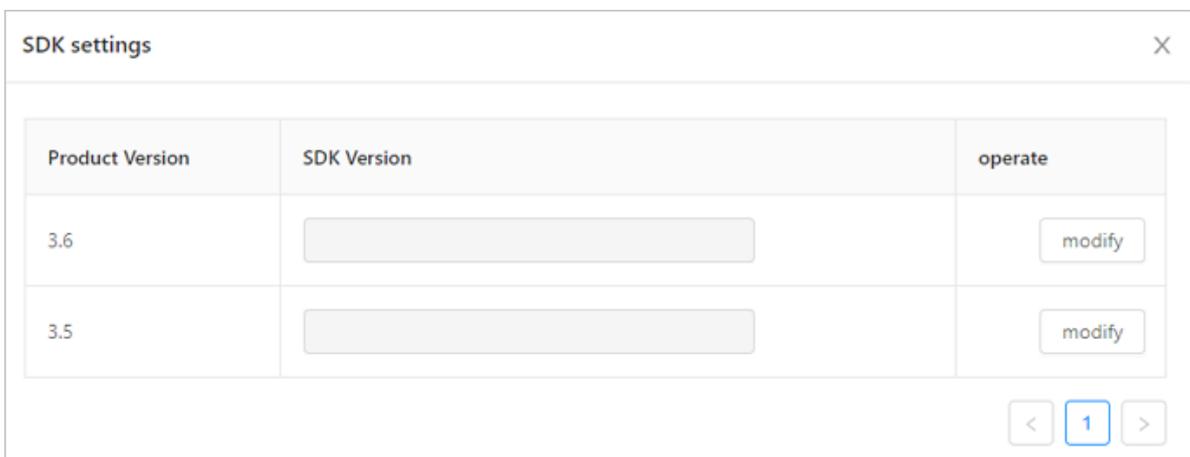
- Click OK to import the API to the system.

1.4.4.3.4 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

- Log on to the Apsara Opsapi Management system.
- In the left-side navigation pane, choose Versions > Products. The Products page appears.
- In the product list, click the  icon in the Actions column corresponding to the product of which the SDK version is to be modified.
- In the SDK Settings dialog box that appears, click Modify in the Actions column corresponding to the product version.



5. Select the specified SDK version from the drop-down list. Click **Submit** in the **Actions** column corresponding to the SDK version. The SDK version is modified.

1.4.4.3.5 Modify product names and descriptions

You can modify the name and description of a product.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. On the **Products** page, click the  icon in the **Actions** column corresponding to the product information about which is to be modified.
4. In the dialog box that appears, modify the product name or description, and click **Submit**.

1.4.4.3.6 View information about product versions

When you need to learn about how to use a product, you can view information about the product version and API version.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. On the **Products** page, click the  icon in the **Actions** column corresponding to the product about which the version information is to be viewed.

You can view information about the product version and API version.

1.4.4.3.7 Modify information about product versions

You can modify information about a product version or API version as needed.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Versions > Products**. The **Products** page appears.
3. On the **Products** page, click the  icon in the **Actions** column corresponding to the product about which the version information is to be modified.

4. In the View Version dialog box that appears, click the  icon in the Actions column corresponding to the version information about which is to be modified.
5. In the dialog box that appears, modify information about the product version and API version.

1.4.4.3.8 Remove information about product versions

You can remove information about a product version that is not applicable.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
3. On the Products page, click the  icon in the Actions column corresponding to the product about which the version information is to be removed.
4. In the View Version dialog box that appears, click the  icon in the Actions column corresponding to the version information about which is to be removed.
5. In the message that appears, click OK.

1.4.4.3.9 Remove information about products

You can remove information about a product that you no longer need.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Versions > Products. The Products page appears.
3. Click the  icon in the Actions column corresponding to the product information about which is to be removed.

1.4.4.3.10 Remove information about product APIs

You can remove information about APIs that are not applicable to a product.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Versions > Products. The Products page appears.

3. In the product list, click the  icon in the Actions column corresponding to the product to be managed.
4. In the dialog box that appears, select the product version from the left drop-down list.

Imported APIs are displayed in the APIs section.
5. Click the  icon corresponding to the API information about which is to be removed.
6. In the message that appears, click Yes.

1.4.4.4 SDK management

The Apsara Opsapi Management system enables you to customize SDKs. You can customize an SDK as needed to export APIs of Apsara Stack products of a specific version. You can also modify and delete the customized SDK.

1.4.4.4.1 Customize SDKs

The Apsara Opsapi Management system provides a tool to customize SDKs. The tool enables you to customize multiple combinations of SDKs for APIs within and across Apsara Stack products of specified versions.

Context

Each product has corresponding SDKs for different programming languages. The Apsara Opsapi Management system supports only Java and Python SDKs.

Each SDK consists of the SDK core and SDK model. The SDK core is the framework of the SDK. It is used to generate HTTP requests or requests of other protocols. The SDK core is fixed. You do not need to generate it each time. The SDK model defines the request parameters and responses of each API.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Versions > SDK Tools. The SDK Tools page appears.
3. In the upper-right corner of the page, click Customize SDK.
4. Set Apsara Stack Version, Product Name, Product Version, SDK Version, API Version, and Language.

The corresponding APIs are displayed in the following APIs section.

5. Select APIs and click Create SDK.

After an SDK is created, you can view the created SDK in the SDK list on the SDK Tools page.

6. Optional: Click the link in the Download column corresponding to the product to download this SDK.

SDK On-demand					
Product Name	Language	Apsara Stack Version	Create Date	Download	Operating
Ecs	Java	v3.8	2019-03-05 08:03:31	ecs-java-sdk_2019-03-05_200330.zip	🔗 🗑️
Ecs	Python	v3.8	2019-03-05 08:02:27	ecs-python-sdk_2019-03-05200226.zip	🔗 🗑️
Bcc	Python	v3.8	2019-03-05 08:01:06	Bcc-python-sdk_2019-03-05_20_01_05.zip	🔗 🗑️
DFS	Java	v3.7	2019-03-05 07:57:51	java-sdk_2019-03-05_19_57_43.zip	🔗 🗑️
opsapi	Java	v3.5	2019-03-05 07:31:55	java-sdk_2019-03-05_19_31_54.zip	🔗 🗑️
opsapi	Python	v3.9	2019-03-05 05:48:05	python-sdk_2019-03-05_17_48_03.zip	🔗 🗑️



Note:

The SDK generated in the Apsara Opsapi Management system is the SDK model. To use this SDK, you need to download the SDK core. You can download the SDK core from the Alibaba Cloud official website or obtain the SDK core from the Apsara Stack after-sales service.

1.4.4.4.2 Modify SDKs

When you need to update an SDK, you can upload an SDK to replace the original SDK.

Procedure

1. [Log on to the Apsara Opsapi Management system.](#)
2. In the left-side navigation pane, choose **Versions > SDK Tools**. The SDK Tools page appears.
3. In the product list, click the [🔗](#) icon in the Actions column corresponding to the product of which the SDK is to be modified.

4. In the dialog box that appears, upload an SDK as prompted and click Submit.

SDK Editor:Bcc-python-sdk_2019-03-05_20_01_05

SDK Version: v3.8

Re-upload SDK:

Click or drag file to this area to upload

submit

1.4.4.4.3 Delete SDKs

You can delete an SDK that you no longer need.

Procedure

1. Log on to the Apsara Opsapi Management system.
2. In the left-side navigation pane, choose Versions > SDK Tools. The SDK Tools page appears.
3. In the product list, click the  icon in the Actions column corresponding to the product of which the SDK is to be deleted.
4. In the message that appears, click Yes.

1.4.5 Test management

To facilitate API tests in the Apsara Opsapi Management system, the system provides the test management function. Each API can be saved as a test case during the test. A test case contains the request parameters of these APIs. You can associate multiple test cases to create a test set. You can choose to run one test case and one test set at a time. You can view execution results on the Execution History page.

1.4.5.1 Test cases

A test case is used to test a specified API.

1.4.5.1.1 Modify test cases

You can modify request parameters of a test case as needed.

Procedure

1. [Log on to the Apsara Opsapi Management system.](#)
2. In the left-side navigation pane, choose **Testing Platform > Test Cases**. The **Test Cases** page appears.
3. On the **Test Cases** page, click the  icon in the **Actions** column corresponding to the test case to be modified.

ID	name	API	product	APIAPI version	description	operation
30	DescribeImages	DescribeImages	Ecs	2014-05-26		 
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, apiId, and apiVersion parameters.

Edit a Test Case [X]

TestCase Name:

Request Parameters

regionId:

accessKeyId:

accessKeySecret:

apiName:

regionNo:

product:

apild:

apiVersion:

5. Click Save to save the modifications.

1.4.5.1.2 Run test cases

You can run a test cases as needed.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Testing Platform > Test Cases**. The **Test Cases** page appears.

3. On the Test Cases page, click the  icon in the Actions column corresponding to the test case.

After the test case is run, view the execution result on the Execution History page.

1.4.5.1.3 Delete test cases

You can delete a test case that you no longer need.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Testing Platform > Test Cases. The Test Cases page appears.
3. On the Test Cases page, click the  icon in the Actions column corresponding to the test case to be deleted.
4. In the message that appears, click Yes.

1.4.5.2 Test sets

A test set consists of multiple associated test cases.

1.4.5.2.1 Create test sets

You can create a test set based on test requirements.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Testing Platform > Test Sets. The Test Sets page appears.
3. On the Test Sets page, click Create Test Set.
4. In the dialog box that appears, enter the test set name and description. Click Save.

We recommend that you enter a test set name that is easily identified.

1.4.5.2.2 Associate test cases

You can associate test cases with a test set to manage test cases in a unified manner.

Procedure

1. *Log on to the Apsara Opsapi Management system.*

2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.
3. On the **Test Sets** page, click the  icon in the **Actions** column corresponding to the test set.
4. **Optional:** You can update the name and description of the test set and click **Save**.
5. Click **Relate to Test Case**.
6. In the dialog box that appears, search for and select the test case to be associated.
You can select multiple test cases and add them to the test set.
7. Click **Save**.

1.4.5.2.3 Run test sets

You can run a test set to check whether the APIs in the test set are available.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.

ID	name	description	operation
25	RDS	test	  
24	ECS	test	  
18	tsc111	desc222	  

3. On the **Test Sets** page, click the  icon in the **Actions** column corresponding to the test case. The test cases in the test set start to run.
After the test set is run, you can view the execution results on the **Execution History** page.

1.4.5.2.4 Delete test sets

You can delete test sets that you no longer need.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **Testing Platform > Test Sets**. The **Test Sets** page appears.

3. On the Test Sets page, click the  icon in the Actions column corresponding to the test case to be deleted.
4. In the message that appears, click Yes.

1.4.6 View execution history of test cases

You can view information about the API for which a test case was executed, including the corresponding product, version information, execution time, and execution status.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose Testing Platform > Execution History. The Execution History page appears.
3. On the Execution History page, click the details icon in the Details column corresponding to the API to be viewed.
4. In the Execution Details dialog box that appears, view the execution details of the test case, including request parameters and responses.

1.4.7 System management

1.4.7.1 Metadatabase management

You can add or remove information about metadatabase in the Apsara Opsapi Management system.

1.4.7.1.1 View information about added metadatabases

The Apsara Opsapi Management system automatically scans all metadatabases in the Apsara Stack environment during initialization. The Apsara Opsapi Management system can scan all metadatabases to view information about the added metadatabases. You can also manually add information about the metadatabases.

Context

The metadatabase information contains the domain name, database name, port, and server that are used in Apsara Stack products.

Procedure

1. *Log on to the Apsara Opsapi Management system.*

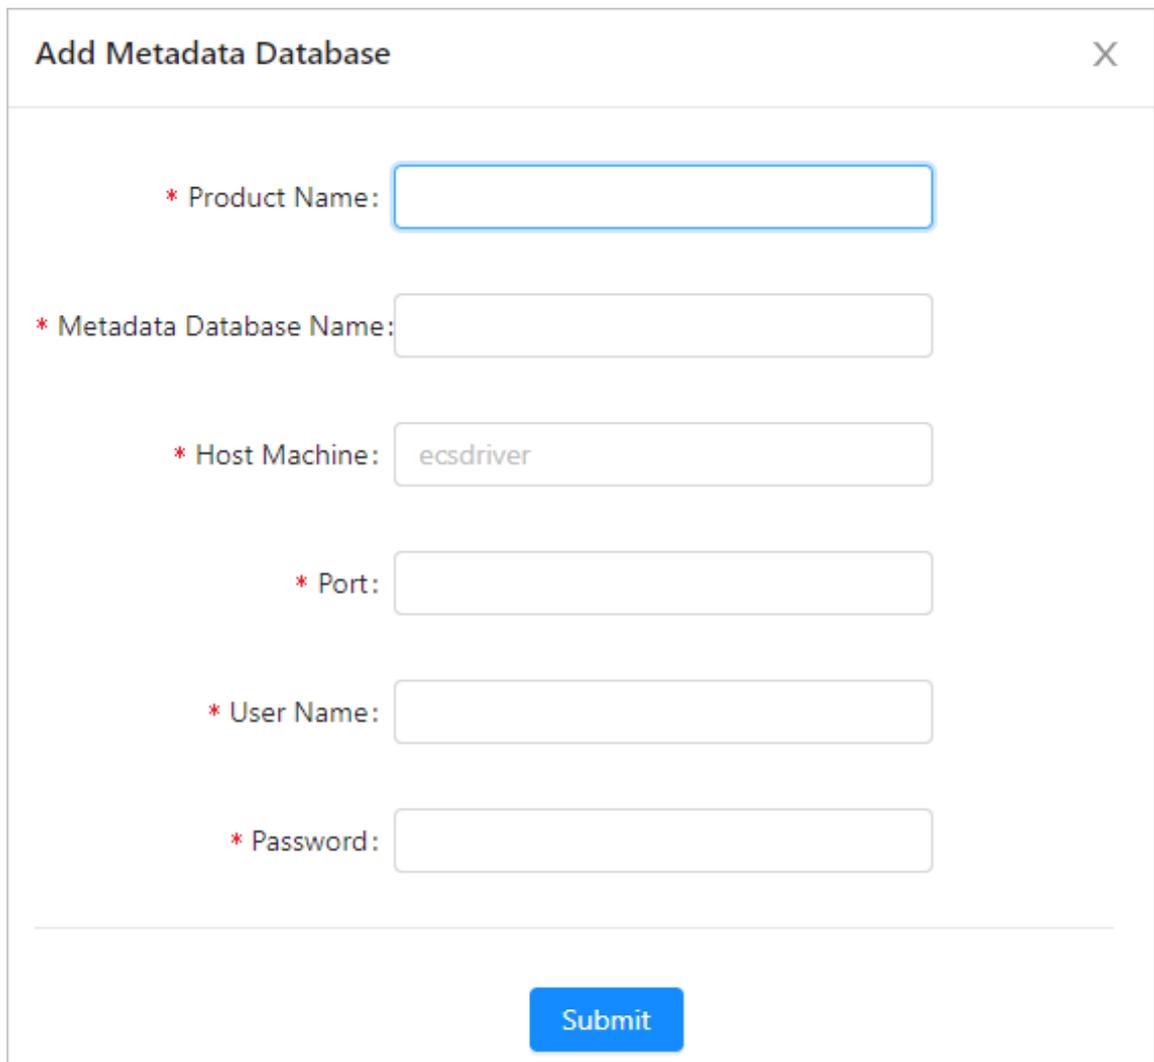
2. In the left-side navigation pane, choose **System Management > Metabase**. The Metabase page appears.
3. Use either of the following methods to view the connection information about metabases:

- **Scan metadatabases**

Click **Scan Metabase** to scan all added metadatabases in the Aspara Stack environment.

- **Add information about metabases**

Click **Add Metabase**. In the **Add Metabase** dialog box that appears, set **Product Name**, **Metabase Name**, **Metabase Server**, **Metabase Port**, **Username**, and **Password**. Click **Submit**.



Add Metadata Database [X]

* Product Name:

* Metadata Database Name:

* Host Machine:

* Port:

* User Name:

* Password:

Submit

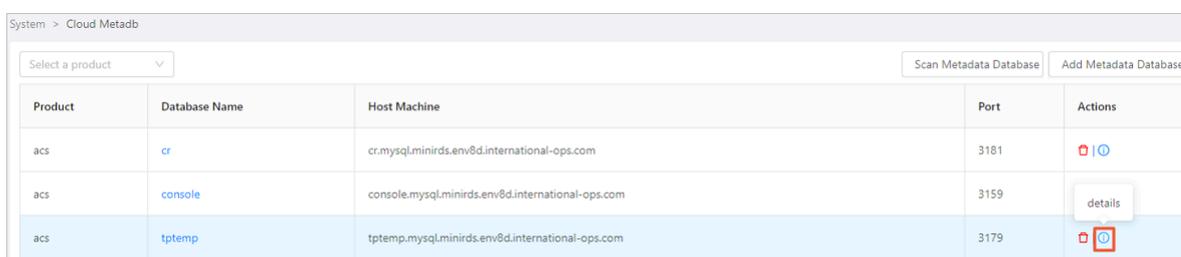
You can view information about the added metadatabases in the metadatabase list.

1.4.7.1.2 View connection information about metadatabases

You can view connection information about a metadatabase on the Metabase page.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose System Management > Metabase. The Metabase page appears.
3. 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.



Product	Database Name	Host Machine	Port	Actions
acs	cr	cr.mysql.minirds.env@d.international-ops.com	3181	
acs	console	console.mysql.minirds.env@d.international-ops.com	3159	
acs	tptemp	tptemp.mysql.minirds.env@d.international-ops.com	3179	 

1.4.7.1.3 Remove information about metadatabases

To facilitate management, you can remove information about metadatabases that you no longer need.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose System Management > Metabase. The Metabase page appears.
3. On the Metabase page, click the  icon in the Actions column corresponding to the metadatabase information about which is to be removed.
4. In the message that appears, click Yes.

1.4.7.2 Server management

You can add or remove information about servers in the Apsara Opsapi Management system.

1.4.7.2.1 View information about added servers

The Apsara Opsapi Management system automatically scans all servers (including physical servers and VMs) in the Apsara Stack environment during initialization.

When new servers are added to the Apsara Stack environment, you can scan servers

to view information about the added servers. You can also add information about the added servers.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. **In the left-side navigation pane, choose System Management > Server Management. The Server Management page appears.**

3. Use either of the following methods to view information about added servers:

- Scan servers

Click Scan Server to scan the information about all servers in the Apsara Stack environment.

- Add information about new servers

Click Add Server. In the Add Server dialog box that appears, set Name, Server Name, Server IP, SSH Port, and SSH User. Upload the SSH private key. Click Submit.

Add host machine ×

* Name:

* Host Name:

* IP Address:

* SSH Port:

* SSH User:

SSH Password:

SSH key:

You can view information about the added servers in the server list.

1.4.7.2.2 Remove server information

To facilitate management, you can remove information about servers that you no longer need.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **System Management > Server Management**. The **Server Management** page appears.
3. On the **Server Management** page, click the  icon in the **Actions** column corresponding to the server information about which is to be removed.

1.4.7.3 Audit APIs

You can view call records of all Opsapis. The records contain the specific API, statuses, time, and result of each call.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **System Management > API Audit**. The **API Audit** page appears.
3. On the **API Audit** page, click the  icon in the **Actions** column corresponding to the API. You can view the call result of this API.

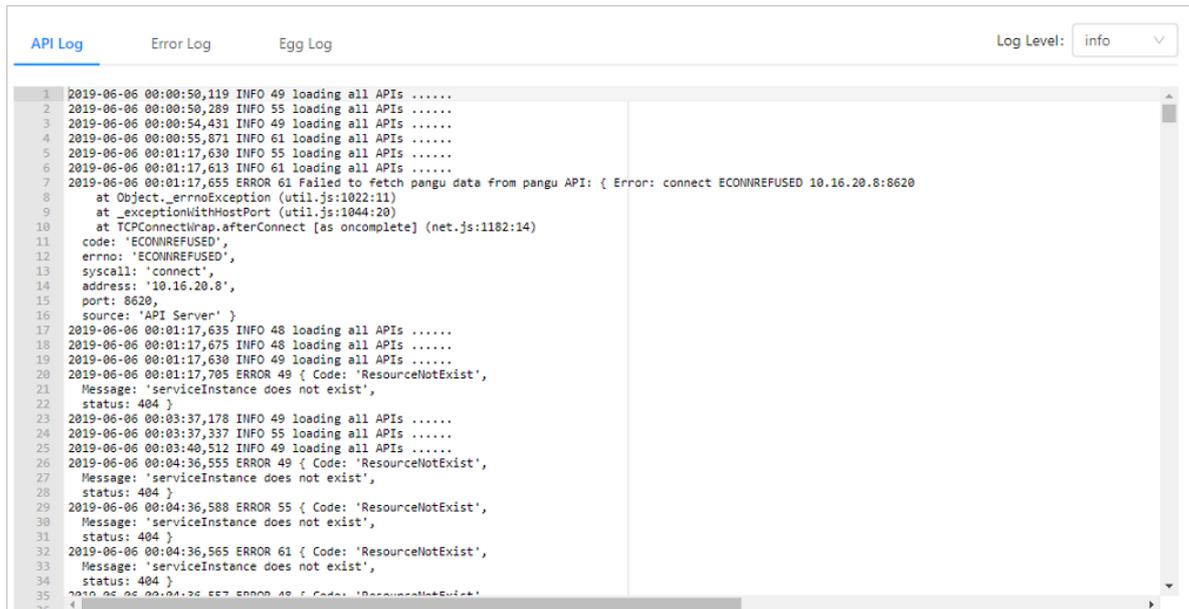
1.4.7.4 View logs

You can view API logs, error logs, and Egg logs to better maintain the back end.

Procedure

1. *Log on to the Apsara Opsapi Management system.*
2. In the left-side navigation pane, choose **System Management > Log Management**. The **Log Management** page appears.

3. View the details of all logs on the API Log, Error Log, and Egg Log tabs.



```

1 2019-06-06 00:00:50,119 INFO 49 loading all APIs .....
2 2019-06-06 00:00:50,289 INFO 55 loading all APIs .....
3 2019-06-06 00:00:54,431 INFO 49 loading all APIs .....
4 2019-06-06 00:00:55,871 INFO 61 loading all APIs .....
5 2019-06-06 00:01:17,630 INFO 55 loading all APIs .....
6 2019-06-06 00:01:17,613 INFO 61 loading all APIs .....
7 2019-06-06 00:01:17,655 ERROR 61 Failed to fetch pangu data from pangu API: { Error: connect ECONNREFUSED 10.16.20.8:8620
8   at Object._errnoException (util.js:1022:11)
9   at _exceptionWithHostPort (util.js:1044:20)
10  at TCPConnectWrap.afterConnect [as oncomplete] (net.js:1182:14)
11  code: 'ECONNREFUSED',
12  errno: 'ECONNREFUSED',
13  syscall: 'connect',
14  address: '10.16.20.8',
15  port: 8620,
16  source: 'API Server' }
17 2019-06-06 00:01:17,635 INFO 48 loading all APIs .....
18 2019-06-06 00:01:17,675 INFO 48 loading all APIs .....
19 2019-06-06 00:01:17,630 INFO 49 loading all APIs .....
20 2019-06-06 00:01:17,705 ERROR 49 { Code: 'ResourceNotExist',
21  Message: 'serviceInstance does not exist',
22  status: 404 }
23 2019-06-06 00:03:37,178 INFO 49 loading all APIs .....
24 2019-06-06 00:03:37,337 INFO 55 loading all APIs .....
25 2019-06-06 00:03:40,512 INFO 49 loading all APIs .....
26 2019-06-06 00:04:36,555 ERROR 49 { Code: 'ResourceNotExist',
27  Message: 'serviceInstance does not exist',
28  status: 404 }
29 2019-06-06 00:04:36,588 ERROR 55 { Code: 'ResourceNotExist',
30  Message: 'serviceInstance does not exist',
31  status: 404 }
32 2019-06-06 00:04:36,565 ERROR 61 { Code: 'ResourceNotExist',
33  Message: 'serviceInstance does not exist',
34  status: 404 }
35 2019-06-06 00:04:36,557 ERROR 49 { Code: 'ResourceNotExist',
36

```



Note:

You can modify log levels in the Apsara Opsapi Management system in real time. The level you set is valid only during the active service operating period. If you restart the service, the default level remains.

1.5 Apsara Infrastructure Management Framework

1.5.1 What is Apsara Infrastructure Management Framework?

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

1.5.1.1 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**

1.5.1.2 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.

1.5.2 Log on to Apsara Infrastructure Management Framework

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

Prerequisites

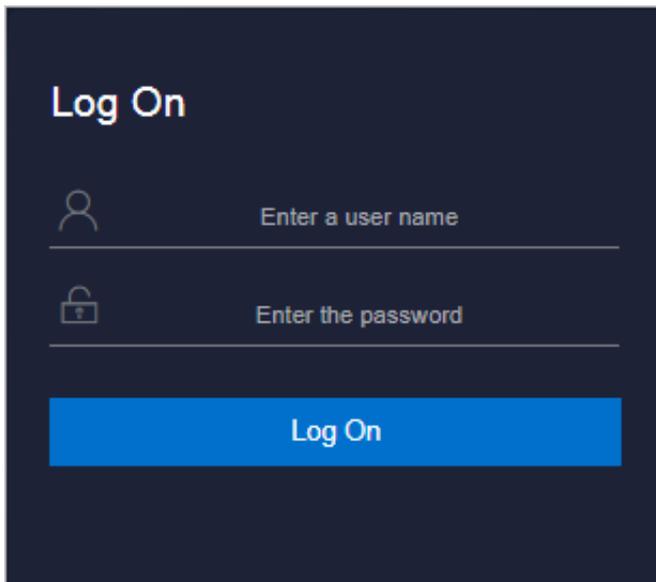
- **ASO access address in the format of `http://region-id.aso.intranet-domain-id.com`.**

- **Google Chrome browser (recommended).**

Procedure

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

Figure 1-9: 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.
5. In the left-side navigation pane, select Products.
6. In the product list, select Apsara Infrastructure Management Framework.

1.5.3 Web page introduction

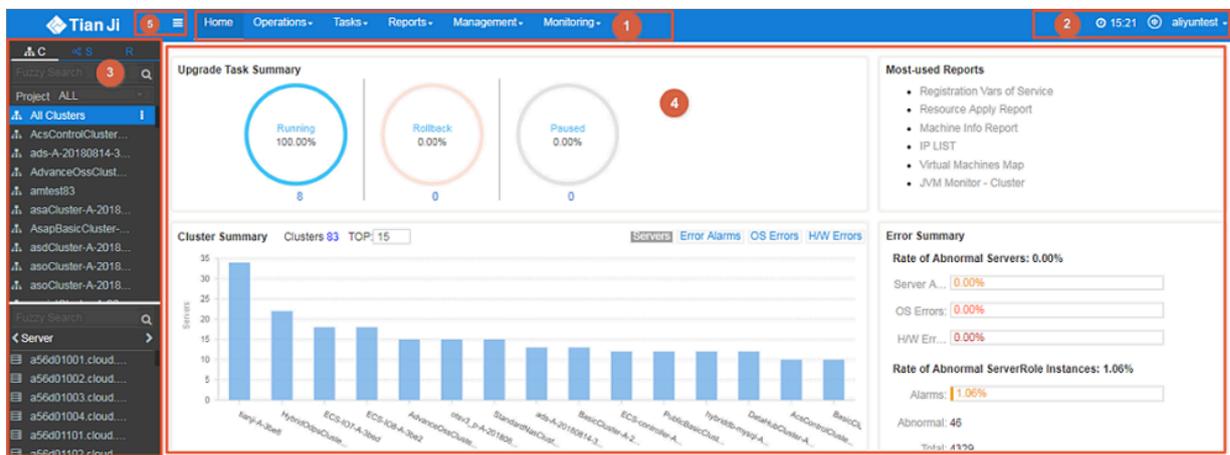
Before performing Operation & Maintenance (O&M) operations on Apsara Infrastructure Management Framework, you must have a general understanding of the Apsara Infrastructure Management Framework page.

1.5.3.1 Introduction on the home page

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

Log on to Apsara Infrastructure Management Framework. The home page appears, as shown in Figure 1-10: Home page of Apsara Infrastructure Management Framework.

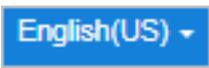
Figure 1-10: Home page of Apsara Infrastructure Management Framework



For more information about the descriptions of functional areas on the home page, see *Table 1-10: Descriptions of functional areas*.

Table 1-10: Descriptions of functional areas

Area		Description
1	Top navigation bar	<ul style="list-style-type: none"> • Operations: the quick entrance of Operation & Maintenance (O&M) operations, which allows operations engineers to quickly find the corresponding operations and operation objects. This menu consists of the following sections: <ul style="list-style-type: none"> - Cluster Operations: performs O&M operations on and manages clusters with the project permissions, such as viewing the cluster status. - Service Operations: manages services with the service permissions, such as viewing the service list information. - Machine Operations: maintains and manages all the machines in Apsara Infrastructure Management Framework, such as viewing the machine status. • Tasks: A rolling task is generated after you modify the configurations in the system. In this menu, you can view running tasks, history tasks, and the deployment summary of clusters, services, and server roles in all projects. • Reports: displays the monitoring data in tables and provides the function of searching for different reports. • Monitoring: effectively monitors metrics in the process of system operation and sends alert notifications for abnormal conditions. This menu includes the functions of displaying alert status, modifying alert rules, and searching for the alert history.

Area		Description
2	Function buttons in the upper-right corner	<ul style="list-style-type: none"> • : - 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. <p>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.</p> <ul style="list-style-type: none"> • : In the English environment, click this drop-down list to switch to another language. • : The logon account information. Click this drop-down list and select Logout to log out of Apsara Infrastructure Management Framework.
3	Left-side navigation pane	<p>In the left-side navigation pane, you can directly view the logical structure of the Apsara Infrastructure Management Framework model.</p> <p>You can view the corresponding detailed data analysis and operations by selecting different levels of nodes in the left-side navigation pane. For more information, see Introduction on the left-side navigation pane.</p>

Area		Description
4	Home page	<p>Displays the summary of related tasks or information as follows:</p> <ul style="list-style-type: none"> • Upgrade Task Summary: the numbers and proportions of running, rolling back, and paused upgrade tasks. • Cluster Summary: the numbers of machines, error alerts, operating system errors, and hardware errors for different clusters. • Error Summary: the metrics for the rate of abnormal machines and the rate of abnormal server role instances. • Most-used Reports: links of the most commonly used statistics reports, which facilitates you to view the report information.
5	Button used to collapse /expand the left-side navigation pane	<p>If you are not required to use the left-side navigation pane when performing O&M operations, click  to collapse the left-side navigation pane and increase the space of the content area.</p>

1.5.3.2 Introduction on the left-side navigation pane

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

Cluster

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

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

- Enter the cluster name in the search box to search for the cluster quickly. Fuzzy search is supported.
- Select a project from the Project drop-down list to display all the clusters in the project.
- Move the pointer over  at the right of a cluster and then perform operations on the cluster as instructed.

- Click a cluster and all the machines and services in this cluster are displayed in the lower-left corner. Move the pointer over  at the right of a machine or service and then perform operations on the machine or service as instructed.
- Click the Machine tab in the lower-left corner. Double-click a machine to view all the server roles in the machine. Double-click a server role to view the applications and then double-click an application to view the log files.
- Click the Service tab in the lower-left corner. Double-click a service to view all the server roles in the service. Double-click a server role to view the machines, double-click a machine to view the applications, and double-click an application to view the log files.
- Double-click a log file. Move the pointer over  at the right of the log file and then select Download to download the log file.

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

Service

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

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

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

Report

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

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

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

1.5.4 Cluster operations

This topic describes the actions about cluster operations.

1.5.4.1 View cluster configurations

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

Procedure

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

The Cluster Operations page displays the following information:

- **Cluster**

The cluster name. Click the cluster name to go to the [Cluster Dashboard](#) page.

- **Scale-Out/Scale-In**

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

- **Abnormal Machine Count**

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

- **Final Status of Normal Machines**

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

- **Rolling**

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

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

For more information about the Cluster Configuration page, see [Table 1-11: Cluster configurations](#).

Table 1-11: Cluster configurations

Category	Item	Description
Basic Information	Cluster	The cluster name.
	Project	The project to which the cluster belongs.
	Clone Switch	<ul style="list-style-type: none"> • Mock Clone: The system is not cloned when a machine is added to the cluster. • Real Clone: The system is cloned when a machine is added to the cluster.
	Machines	The number of machines in the cluster. Click View Clustering Machines to view the machine list.
	Security Verification	The access control among processes. Generally, the non-production environment uses the default configurations and does not perform the verification. In other cases, customize the configurations based on actual requirements to enable or disable the verification.
	Cluster Type	<ul style="list-style-type: none"> • RDS • NETFRAME • T4: a special type that is required by the mixed deployment of e-commerce. • Default: other conditions.
Deployment Plan	Service	The service deployed in the cluster.

Category	Item	Description
	Dependency Service	The service that the current service depends on.
Service Information	Service Information	Select a service from the Service Information drop-down list and then the configurations of this service are displayed.
	Service Template	The template used by the service.
	Monitoring Template	The monitoring template used by the service.
	Machine Mappings	The machines included in the server role of the service.
	Software Version	The software version of the server role in the service.
	Availability Configuration	The availability configuration percentage of the server role in the service.
	Deployment Plan	The deployment plan of the server role in the service.
	Configuration Information	The configuration file used in the service.
	Role Attribute	Server roles and the corresponding parameters.

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

1.5.4.2 View the cluster dashboard

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

Procedure

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

2. You have two ways to go to the Cluster Dashboard page:

- In the left-side navigation pane, click the C tab. Move the pointer over  at the right of a cluster and then select Dashboard.
- In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, click the cluster name.

3. On the Cluster Dashboard page, you can view the cluster information, including the basic information, final status information, rolling job information, dependencies, resource information, virtual machines, and monitoring information. For more information about the descriptions, see the following table.

Item	Description
Basic Cluster Information	<p>Displays the basic information of the cluster as follows:</p> <ul style="list-style-type: none"> • Project Name: the project name. • Cluster Name: the cluster name. • IDC: the data center to which the cluster belongs. • Final Status Version: the latest version of the cluster. • Cluster in Final Status: whether the cluster reaches the final status. • Machines Not In Final Status: the number of machines that do not reach the final status in the cluster when the cluster does not reach the final status. • Real/Pseudo Clone: whether to clone the system when a machine is added to the cluster. • Expected Machines: the number of expected machines in the cluster. • Actual Machines: the number of machines in the current environment. • Machines Not Good: the number of machines whose status is not Good in the cluster. • Actual Services: the number of services that are actually deployed in the cluster. • Actual Server Roles: the number of server roles that are actually deployed in the cluster. • Cluster Status: whether the cluster is starting or shutting down machines.

Item	Description
Machine Status Overview	The statistical chart of the machine status in the cluster .
Machines in Final Status	The numbers of machines that reach the final status and those that do not reach the final status in each service of the cluster.
Load-System	The system load chart of the cluster.
CPU-System	The CPU load chart.
Mem-System	The memory load chart.
Disk_usage-System	The statistical table of the disk usage.
Traffic-System	The system traffic chart.
TCP State-system	The TCP request status chart.
TCP Retrans-System	The chart of TCP retransmission amount.
Disk_IO-System	The statistical table of the disk input and output.
Service Instances	<p>Displays the service instances deployed in the cluster and the related final status information.</p> <ul style="list-style-type: none"> • Service Instance: the service instance deployed in the cluster. • Final Status: whether the service instance reaches the final status. • Expected Server Roles: the number of server roles that the service instance expects to deploy. • Server Roles In Final Status: the number of server roles that reach the final status in the service instance. • Server Roles Going Offline: the number of server roles that are going offline in the service instance. • Actions: Click Details to go to the Service Instance Information Dashboard page. For more information about the service instance dashboard, see View the service instance dashboard.

Item	Description
Upgrade Tasks	<p>Displays the upgrade tasks related to the cluster.</p> <ul style="list-style-type: none"> • Cluster Name: the name of the upgrade cluster. • Type: the type of the upgrade task. The options include app (version upgrade) and config (configuration change). • Git Version: the change version to which the upgrade task belongs. • Description: the description about the change. • Rolling Result: the result of the upgrade task. • Submitted By: the person who submits the change. • Submitted At: the time when the change is submitted. • Start Time: the time to start the rolling. • End Time: the time to finish the upgrade. • Time Used: the time used for the upgrade. • Actions: Click Details to go to the Rolling Task page. For more information about the rolling task, see View rolling tasks.
Cluster Resource Request Status	<ul style="list-style-type: none"> • Version: the resource request version. • Msg: the exception message. • Begintime: the start time of the resource request analysis. • Endtime: the end time of the resource request analysis. • Build Status: the build status of resources. • Resource Process Status: the resource request status in the version.

Item	Description
Cluster Resource	<ul style="list-style-type: none"> • Service: the service name. • Server Role: the server role name. • App: the application of the server role. • Name: the resource name. • Type: the resource type. • Status: the resource request status. • Error Msg: the exception message. • Parameters: the resource parameters. • Result: the resource request result. • Res: the resource ID. • Reprocess Status: the status of interaction with Business Foundation System during the VIP resource request. • Reprocess Msg: the exception message of interaction with Business Foundation System during the VIP resource request. • Reprocess Result: the result of interaction with Business Foundation System during the VIP resource request. • Refer Version List: the version that uses the resource.
VM Mappings	<p>The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.</p> <ul style="list-style-type: none"> • VM: the hostname of the virtual machine. • Currently Deployed On: the hostname of the physical machine where the virtual machine is currently deployed. • Target Deployed On: the hostname of the physical machine where the virtual machine is expected to be deployed.

Item	Description
Service Dependencies	<p>The dependencies of service instances and server roles in the cluster, and the final status information of the dependent service or server role.</p> <ul style="list-style-type: none"> • Service: the service name. • Server Role: the server role name. • Dependent Service: the service on which the server role depends. • Dependent Server Role: the server role on which the server role depends. • Dependent Cluster: the cluster to which the dependent server role belongs. • Dependency in Final Status: whether the dependent server role reaches the final status.

1.5.4.3 View the cluster operation and maintenance center

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

Procedure

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

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

Item	Description
SR not in Final Status	<p>Displays all the server roles that do not reach the final status in the cluster.</p> <p>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.</p>
Running Tasks	<p>Displays whether the cluster has running rolling tasks.</p> <p>Click Rolling to go to the Rolling Task page. For more information about the rolling task, see View rolling tasks.</p>
Head Version Submitted At	<p>The time when the head version is submitted.</p> <p>Click the time to view the submission details.</p>
Head Version Analysis	<p>The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to detailed change contents. The head version analysis has the following statuses:</p> <ul style="list-style-type: none"> • Preparing: No new version is available now. • Waiting: The latest version is found. The analysis module has not started up yet. • Doing: The module is analyzing the application that requires change. • done: The head version analysis is successfully completed. • Failed: The head version analysis failed. The change contents cannot be parsed. <p>If the status is not done, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version.</p> <p>Click the status to view the relevant information.</p>

Item	Description
Service	Select a service deployed in the cluster from the drop-down list.
Server Role	<p>Select a server role of a service in the cluster from the drop-down list.</p> <div data-bbox="571 465 1436 672" style="background-color: #f0f0f0; padding: 5px;">  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. </div>
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	<p>The number of abnormal machines that encounter each type of the following faults.</p> <ul style="list-style-type: none"> • Ping Failed: A ping_monitor error is reported, and TianjiMaster cannot successfully ping the machine. • No Heartbeat: TianjiClient on the machine does not regularly report data to indicate the status of this machine, which may be caused by the TianjiClient problem or network problem. • Status Error: The machine has an error reported by the monitor or a fault of the critical or fatal level. Check the alert information and accordingly solve the issue.

Item	Description
Abnormal Services	<p>The number of machines with abnormal services. To determine if a service reaches the final status, see the following rules:</p> <ul style="list-style-type: none">• The server role on the machine is in the GOOD status.• Each application of the server role on the machine must keep the actual version the same as the head version.• Before the Image Builder builds an application of the head version, Apsara Infrastructure Management Framework cannot determine the value of the head version and the service final status is unknown. This process is called the change preparation process. The service final status cannot be determined during the preparation process or upon a preparation failure.

Item	Description
Machines	<p>Displays all the machines in the cluster or the machines included in a specific server role of a specific service.</p> <ul style="list-style-type: none"> • Machine search: Click the search box to enter the machine in the displayed dialog box. Fuzzy or batch search is supported. • Click the machine name to view the physical information of the machine in the displayed Machine Information dialog box. Click Dashboard to go to the Machine Details page. For more information about the machine details, see View the machine dashboard. • Move the pointer over the blank area in the Final Status column or the Final SR Status column and then click Details to view the machine status, system service information, server role status on the machine, and exception message. • If no service or server role is selected from the drop-down list, move the pointer over the blank area in the Running Status column and then click Details to view the running status information or exception message of the machine. <p>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.</p> <ul style="list-style-type: none"> • Click Error, Warning, or Good in the Monitoring Statistics column to view the monitored items of machines and monitored items of server roles. • Click Terminal in the Actions column to log on to the machine and perform related operations. • Click Machine Operation in the Actions column to restart, out-of-band restart, or clone the machine again.

1.5.4.4 View the service final status

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

Procedure

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

2. You have two ways to go to the Service Final Status Query page:

- In the left-side navigation pane, click the C tab. Move the pointer over  at the right of a cluster and then choose Monitoring > Service Final Status Query.
- In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, choose Monitoring > Service Final Status Query in the Actions column at the right of a cluster.

3. View the information on the Service Final Status Query page.

Item	Description
Project Name	The name of the project to which the cluster belongs.
Cluster Name	The cluster name.
Head Version Submitted At	The time when the head version is submitted.
Head Version Analysis	<p>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:</p> <ul style="list-style-type: none"> • Preparing: No new version is available now. • Waiting: The latest version is found. The analysis module has not started up yet. • Doing: The module is analyzing the application that requires change. • done: The head version analysis is successfully completed. • Failed: The head version analysis failed. The change contents cannot be parsed. <p>If the status is not done, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version.</p>
Cluster Rolling Status	Displays the information of the current rolling task in the cluster, if any. The rolling task may not be of the head version.

Item	Description
Cluster Machine Final Status Statistics	The status of all machines in the cluster. Click View Details to go to the Cluster Operation and Maintenance Center page and view the detailed information of all machines. For more information about the cluster operation and maintenance center, see View the cluster operation and maintenance center .
Final Status of Cluster SR Version	The final status of cluster service version.  Note: Take statistics of services that do not reach the final status, which is caused by version inconsistency or status exceptions. If services do not reach the final status because of machine problems, go to Cluster Machine Final Status Statistics to view the statistics.
Final Status of SR Version	The number of machines that do not reach the final status when a server role has tasks.

1.5.4.5 View operation logs

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. You have two ways to go to the Cluster Operation Logs page:
 - In the left-side navigation pane, click the C tab. Move the pointer over  at the right of a cluster and then choose Monitoring > Operation Logs.
 - In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page, choose Monitoring > Operation Logs in the Actions column at the right of a cluster.
3. On the Cluster Operation Logs page, click Refresh. View the Git version, description, submitter, submitted time, and task status.

4. **Optional: Complete the following steps to view the differences between versions on the Cluster Operation Logs page.**

a) Find the log in the operation log list and then click **View Release Changes** in the **Actions** column.

b) On the **Version Difference** page, complete the following configurations:

- **Select Base Version:** Select a base version.
- **Configuration Type:** Select **Extended Configuration** or **Cluster Configuration**. **Extended Configuration** displays the configuration differences after the configuration on the cluster is combined with the configuration in the template. **Cluster Configuration** displays the configuration differences on the cluster.

c) Click **Obtain Difference**.

The differential file list is displayed.

d) Click each differential file to view the detailed differences.

1.5.5 Service operations

This topic describes the actions about service operations.

1.5.5.1 View the service list

The service list allows you to view the list of all services and the related information.

Procedure

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

Item	Description
Service	The service name.
Service Instances	The number of service instances in the service.
Service Configuration Templates	The number of service configuration templates.
Monitoring Templates	The number of monitoring templates.

Item	Description
Service Schemas	The number of service configuration validation templates.
Actions	Click Management to view the service instances, service templates, monitoring templates, monitoring instances, service schemas, and detection scripts.

1.5.5.2 View the service instance dashboard

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

Procedure

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

6. View the information on the Service Instance Information Dashboard page.

Item	Description
Service Instance Summary	<p>Displays the basic information of the service instance as follows:</p> <ul style="list-style-type: none"> • Cluster Name: the name of the cluster to which the service instance belongs. • Service Name: the name of the service to which the service instance belongs. • Actual Machines: the number of machines in the current environment. • Expected Machines: the number of machines that the service instance expects. • Target Total Server Roles: the number of server roles that the service instance expects. • Actual Server Roles: the number of server roles in the current environment. • Template Name: the name of the service template used by the service instance. • Template Version: the version of the service template used by the service instance. • Schema: the name of the service schema used by the service instance. • Monitoring System Template: the name of the monitoring system template used by the service instance.
Server Role Statuses	The statistical chart of the current status of server roles in the service instance.
Machine Statuses for Server Roles	The status statistics of machines where server roles are located.
Service Monitoring Information	<ul style="list-style-type: none"> • Monitored Item: the name of the monitored item. • Level: the level of the monitored item. • Description: the description of the monitored contents. • Updated At: the time when the data is updated.

Item	Description
Service Alert Status	<ul style="list-style-type: none"> • Alert Name • Instance Information • Alert Start • Alert End • Alert Duration • Severity Level • Occurrences: the number of times the alert is triggered.
Server Role List	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • Alert Name • Alert Time • Instance Information • Severity Level • Contact Group
Service Dependencies	<p>The dependencies of service instances and server roles in the service instance, and the final status information of the dependent service or server role.</p> <ul style="list-style-type: none"> • Server Role: the server role name. • Dependent Service: the service on which the server role depends. • Dependent Server Role: the server role on which the server role depends. • Dependent Cluster: the cluster to which the dependent server role belongs. • Dependency in Final Status: whether the dependent server role reaches the final status.

1.5.5.3 View the server role dashboard

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

Procedure

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

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

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

Item	Description
Machine Information	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> - Click Terminal to log on to the machine and perform operations. - Click Restart to restart the server roles on the machine. - Click Details to go to the Machine Details page. For more information about the machine details, see View the machine dashboard. - Click Machine System View to go to the Machine Info Report page. For more information about the machine info report, see Machine info report. - Click Machine Operation to restart, out of band restart, or clone the machine again.
Server Role Monitoring Information of Machines	<ul style="list-style-type: none"> • Updated At: the time when the data is updated. • Machine Name: the machine name. • Monitored Item: the name of the monitored item. • Level: the level of the monitored item. • Description: the description of the monitored item.

Item	Description
VM Mappings	<p>The information of virtual machines in the cluster. Data is available only when virtual machines are deployed in the cluster.</p> <ul style="list-style-type: none"> • VM: the hostname of the virtual machine. • Currently Deployed On: the hostname of the physical machine where the virtual machine is currently deployed. • Target Deployed On: the hostname of the physical machine where the virtual machine is expected to be deployed.
Service Dependencies	<p>The dependencies of service instances and server roles in the service instance, and the final status information of the dependent service or server role.</p> <ul style="list-style-type: none"> • Dependent Service: the service on which the server role depends. • Dependent Server Role: the server role on which the server role depends. • Dependent Cluster: the cluster to which the dependent server role belongs. • Dependency in Final Status: whether the dependent server role reaches the final status.

1.5.6 Machine operations

This topic describes the actions about machine operations.

1.5.6.1 View the machine dashboard

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the left-side navigation pane, click the C tab.
3. On the Machine tab in the lower-left corner, enter the machine name in the search box. Machines that meet the search condition are displayed.
4. Move the pointer over  at the right of a machine and then select Dashboard.

5. On the Machine Details page, view all the information of this machine. For more information, see the following table.

Item	Description
Load-System	The system load chart of the cluster.
CPU-System	The CPU load chart.
Mem-System	The memory load chart.
DISK Usage-System	The statistical table of the disk usage.
Traffic-System	The system traffic chart.
TCP State-System	The TCP request status chart.
TCP Retrans-System	The chart of TCP retransmission amount.
DISK IO-System	The statistical table of the disk input and output.
Machine Summary	<ul style="list-style-type: none"> • Project Name: the name of the project to which the machine belongs. • Cluster Name: the name of the cluster to which the machine belongs. • Machine Name: the machine name. • SN: the serial number of the machine. • IP: the IP address of the machine. • IDC: the data center of the machine. • Room: the room in the data center where the machine is located. • Rack: the rack where the machine is located. • Unit in Rack: the location of the rack. • Warranty: the warranty of the machine. • Purchase Date: the date when the machine is purchased. • Machine Status: the running status of the machine. • Status: the hardware status of the machine. • CPUs: the number of CPUs for the machine. • Disks: the disk size. • Memory: the memory size. • Manufacturer: the machine manufacturer. • Model: the machine model. • os: the operating system of the machine. • part: the disk partition.
Server Role Status of Machine	The distribution of the current status of all server roles on the machine.

Item	Description
Machine Monitoring Information	<ul style="list-style-type: none"> • Monitored Item: the name of the monitored item. • Level: the level of the monitored item. • Description: the description of the monitored contents. • Updated At: the time when the monitoring information is updated.
Machine Server Role Status	<ul style="list-style-type: none"> • Service Instance • Server Role • Server Role Status • Server Role Action • Error Message • Target Version • Current Version • Actual Version Update Time • Actions: <ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> • Application Name: the application name. • Process Number • Status: the application status. • Current Build ID: the ID of the current package version. • Target Build ID: the ID of the expected package version. • Git Version • Start Time • End Time • Interval: the interval between the time when Apsara Infrastructure Management Framework detects that the process exits and the time when Apsara Infrastructure Management Framework repairs the process. • Information Message: the normal output logs. • Error Message: the abnormal logs.

1.5.7 Monitoring center

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

1.5.7.1 Modify an alert rule

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the top navigation bar, choose **Operations > Service Operations**.
3. Enter the service name in the search box.
4. Find the service and then click **Management** in the **Actions** column.
5. Click the **Monitoring Template** tab.
6. Find the monitoring template that you are about to edit and then click **Edit** in the **Actions** column.
7. Configure the monitoring parameters based on actual conditions.
8. Click **Save Change**.

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

1.5.7.2 View the status of a monitoring instance

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

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. In the top navigation bar, choose **Operations > Service Operations**.
3. Enter the service name in the search box.
4. Find the service and then click **Management** in the **Actions** column.
5. Click the **Monitoring Instance** tab.

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

1.5.7.3 View the alert status

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

Procedure

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

Item	Description
Service	The service name.
Cluster	The name of the cluster where the service is located.
Instance	The name of the service instance being monitored. Click the instance to view the alert history of this instance.
Alert Status	Alerts have two statuses: Restored and Alerting.
Alert Level	Alerts have the following four levels, from high to low, according to the effect on services. <ul style="list-style-type: none"> • P1 • P2 • P3 • P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered and how long the alert has lasted.
Actions	Click Show to show the data before and after the alert time.

1.5.7.4 View alert rules

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

Procedure

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

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

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

1.5.7.5 View the alert history

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

Procedure

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

Item	Description
Service	The name of the service to which the alert belongs.
Cluster	The name of the cluster where the service is located.
Alert Instance	The name of the resource where the alert is triggered.
Status	Alerts have two statuses: Restored and Alerting .

Item	Description
Alert Level	Alerts have the following four levels, from high to low, according to the effect on services. <ul style="list-style-type: none"> • P1 • P2 • P3 • P4
Alert Name	The name of the generated alert. Click the alert name to view the alert rule details.
Alert Time	The time when the alert is triggered.
Alert Contact	The groups and members that are notified when an alert is triggered.
Actions	Click Show to show the data before and after the alert time.

1.5.8 Tasks and deployment summary

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

1.5.8.1 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.

Table 1-12: Change Task list

Item	Description
Change Version	The version that triggers the change of the rolling task.
Description	The description about the change.

Item	Description
Head Version Analysis	<p>The head version analysis is the process that Apsara Infrastructure Management Framework detects the latest cluster version and parses the version to detailed change contents. The head version analysis has the following statuses:</p> <ul style="list-style-type: none"> • Preparing: No new version is available now. • Waiting: The latest version is found. The analysis module has not started up yet. • Doing: The module is analyzing the application that requires change. • done: The head version analysis is successfully completed. • Failed: The head version analysis failed. The change contents cannot be parsed. <p>If the status is not done, Apsara Infrastructure Management Framework cannot detect the change contents of server roles in the latest version.</p>
Blocked Server Role	Server roles blocked in the rolling task. Generally, server roles are blocked because of dependencies.
Submitter	The person who submits the change.
Submitted At	The time when the change is submitted.
Actions	<p>Click View Difference to go to the Version Difference page. For more information, see View operation logs.</p> <p>Click Stop to stop the rolling task.</p> <p>Click Pause to pause the rolling task.</p>

Table 1-13: Change Details list

Item	Description
Service Name	The name of the service where a change occurs.

Item	Description
Status	<p>The current status of the service.</p> <p>The rolling status of the service is an aggregated result, which is calculated based on the rolling status of the server role.</p> <ul style="list-style-type: none"> • succeeded: The task is successfully run. • blocked: The task is blocked. • failed: The task failed.
Server Role Status	<p>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.</p> <p>Server roles have the following statuses:</p> <ul style="list-style-type: none"> • Downloading: The task is being downloaded. • Rolling: The rolling task is running. • RollingBack: The rolling task failed and is rolling back.
Depend On	<p>The services that this service depends on or server roles that this server role depends on.</p>
Actions	<p>Click Stop to stop the change of the server role.</p> <p>Click Pause to pause the change of the server role.</p>

1.5.8.2 View running tasks

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

Procedure

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

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

1.5.8.3 View history tasks

You can view the historical running conditions of completed tasks.

Procedure

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

1.5.8.4 View the deployment summary

On the **Deployment Summary** page, you can view the deployment conditions of clusters, services, and server roles in all projects on Apsara Infrastructure Management Framework.

Procedure

1. [Log on to Apsara Infrastructure Management Framework](#).
2. In the top navigation bar, choose **Tasks > Deployment Summary**.
 - **View the deployment status and the duration of a certain status for each project.**
 - **Gray: wait to be deployed.** It indicates that some services of the project depend on server roles or service instances that are being deployed, and

other service instances or server roles in the project have already been deployed.

- **Blue: being deployed.** It indicates that the project has not reached the final status for one time yet.
- **Green: has reached the final status.** It indicates that all clusters in the project have reached the final status.
- **Orange: not reaches the final status.** It indicates that a server role does not reach the final status for some reason after the project reaches the final status for the first time.
- **Configure the global clone switch.**
 - **normal: Clone is allowed.**
 - **block: Clone is forbidden.**
- **Configure the global dependency switch.**
 - **normal: All configured dependencies are checked.**
 - **ignore: The dependency is not checked.**
 - **ignore_service: None of the service-level dependencies, including the server role dependencies across services, are checked, and only the server role-level dependencies are checked.**

3. Click the Deployment Details tab to view the deployment details.

For more information, see the following table.

Item	Description
Status Statistics	<p>The general statistics of deployment conditions, including the total number of projects that are currently available. Click each status to display the projects in the corresponding status in the list. The projects have five deployment statuses:</p> <ul style="list-style-type: none"> • Final: All the clusters in the project have reached the final status. • Deploying: The project has not reached the final status for one time yet. • Waiting: Some services of the project depend on server roles or service instances that are being deployed, and other service instances or server roles in the project have already been deployed. • Non-final: A server role does not reach the final status for some reason after the project reaches the final status for the first time. • Inspector Warning: An error is detected on service instances in the project during the inspection.
Start Time	The time when Apsara Infrastructure Management Framework starts the deployment.
Progress	The proportion of server roles that reach the final status to all the server roles in the current environment.
Deployment Status	<p>The time indicates the deployment duration for the following statuses: Final, Deploying, Waiting, and Inspector Warning.</p> <p>The time indicates the duration before the final status is reached for the Non-final status.</p> <p>Click the time to view the details.</p>

Item	Description
Deployment Progress	<p>The proportion of clusters, services, and server roles that reach the final status to the total clusters, services, and server roles in the project.</p> <p>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.</p>
Resource Application Progress	<p>Total indicates the total number of resources related to the project.</p> <ul style="list-style-type: none"> • Done: the number of resources that have been successfully applied for. • Doing: the number of resources that are being applied for 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.

1.5.9 Reports

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

1.5.9.1 View reports

The Reports menu allows you to view the statistical data.

Context

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

- System reports: default and common reports in the system.

- **All reports:** includes the system reports and custom reports.

Procedure

1. [Log on to Apsara Infrastructure Management Framework.](#)
2. You can go to the report list in the following three ways:
 - In the top navigation bar, choose Reports > System Reports.
 - In the top navigation bar, choose Reports > All Reports.
 - In the left-side navigation pane, click the R tab. Move the pointer over  at the right of All Reports and then select View.

See the following table for the report descriptions.

Item	Description
Report	The report name. Move the pointer over  next to Report to search for reports by report name.
Group	The group to which the report belongs. Move the pointer over  next to Group to filter reports by group name.
Status	Indicates whether the report is published.
Public	Indicates whether the report is public.
Created By	The person who creates the report.
Published At	The published time and created time of the report.
Actions	Click Add to Favorites to add this report to your favorites. Then, you can view the report by choosing Reports > Favorites in the top 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](#).

1.5.9.2 Add a report to favorites

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

Procedure

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

1.5.10 Appendix

1.5.10.1 Project component info report

This report displays the name and status for each type of project components, including services, server roles, and machines.

Item	Description
Project	The project name.
Cluster	The name of a cluster in the project.
Service	The name of a service in the cluster.
Server Role	The name of a server role in the service.
Server Role Status	The running status of the server role on the machine.
Server Role Action	The action that the server role performs on the machine. Data is available only when Apsara Infrastructure Management Framework asks the server role to perform certain actions, such as rolling and restart actions.
Machine Name	The hostname of the machine.

Item	Description
IP	The IP address of the machine.
Machine Status	The running status of the machine.
Machine Action	The action that Apsara Infrastructure Management Framework asks the machine to perform, such as the clone action.

1.5.10.2 IP list

This report displays the IP addresses of physical machines and Docker applications.

IP List of Physical Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The hostname of the machine.
IP	The IP address of the machine.

IP List of Docker Applications

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
Machine Name	The hostname of the machine.
Docker Host	The Docker hostname.
Docker IP	The Docker IP address.

1.5.10.3 Machine info report

This report displays the statuses of machines and server roles on the machines.

Machine Status

Displays all the machines currently managed by Apsara Infrastructure Management Framework and their corresponding statuses. In the Global Filter section at the top of the page, select the project, cluster, and machine from the project, cluster, and machine drop-down lists, and then click Filter on the right to filter the data.

Item	Description
Machine Name	The machine name.
IP	The IP address of the machine.
Machine Status	The machine status.
Machine Action	The action currently performed by the machine.
Machine Action Status	The action status.
Status Description	The description about the machine status.

Expected Server Role List

Select a row in the Machine Status section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Server Role	The name of the expected server role on the machine.

Abnormal Monitoring Status

Select a row in the Machine Status section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.
Monitored Item	The name of the monitored item.
Level	The level of the monitored item.
Description	The description of the monitored item contents.
Updated At	The updated time of the monitored item.

Server Role Version and Status on Machine

Select a row in the Machine Status section to display the corresponding information in this list.

Item	Description
Machine Name	The machine name.

Item	Description
Server Role	The server role name.
Server Role Status	The status of the server role.
Target Version	The expected version of the server role on the machine.
Current Version	The current version of the server role on the machine.
Status Description	The description about the status.
Error Message	The exception message of the server role.

Monitoring Status

Select a row in the Machine Status section to display the corresponding information in this list.

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

1.5.10.4 Rolling info report

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

Choose a rolling action

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

Item	Description
Cluster	The cluster name.
Git Version	The version of change that triggers the rolling task.
Description	The description about the change entered by a user when the user submits the change.

Item	Description
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Submitted By	The ID of the user who submits the change.
Rolling Task Status	The current status of the rolling task.
Submitted At	The time when the change is submitted.

Server Role in Job

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

Item	Description
Server Role	The server role name.
Server Role Status	The rolling status of the server role.
Error Message	The exception message of the rolling task.
Git Version	The version of change to which the rolling task belongs.
Start Time	The start time of the rolling task.
End Time	The end time of the rolling task.
Approve Rate	The proportion of machines that have the rolling task approved by the decider.
Failure Rate	The proportion of machines that have the rolling task failed.
Success Rate	The proportion of machines that have the rolling task succeeded.

Server Role Rolling Build Information

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

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

Server Role Statuses on Machines

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

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

1.5.10.5 Machine RMA approval pending list

Some Apsara Infrastructure Management Framework actions (such as restart) on machines and server roles can be triggered by users, but this type of actions must be reviewed and approved. This report is used to process the actions that must be reviewed and approved.

Machine

Displays the basic information of pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.

Item	Description
IP	The IP address of the machine.
State	The running status of the machine.
Action Name	The action on the machine.
Action Status	The status of the action on the machine.
Actions	The approval button.

Machine Serverrole

Displays the information of server roles on the pending approval machines.

Item	Description
Project	The project name.
Cluster	The cluster name.
Hostname	The hostname of the machine.
IP	The IP address of the machine.
Serverrole	The server role name.
State	The running status of the server role.
Action Name	The action on the server role.
Action Status	The status of the action on the server role.
Actions	The approval button.

Machine Component

Displays the hard disk information of pending approval machines.

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

1.5.10.6 Registration vars of services

This report displays values of all service registration variables.

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

1.5.10.7 Virtual machine mappings

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

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

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

1.5.10.8 Service inspector report

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

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

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

1.5.10.9 Resource application report

In the Global Filter section, select the project, cluster, and machine from the project, cluster, and machine drop-down lists and then click Filter on the right to display the corresponding resource application data.

Change Mappings

Item	Description
Project	The project name.
Cluster	The cluster name.
Version	The version where the change occurs.
Resource Process Status	The resource application status in the version.
Msg	The exception message.
Begintime	The start time of the change analysis.
Endtime	The end time of the change analysis.

Changed Resource List

Item	Description
Res	The resource ID.
Type	The resource type.
Name	The resource name.
Owner	The application to which the resource belongs.
Parameters	The resource parameters.
Ins	The resource instance name.
Instance ID	The resource instance ID.

Resource Status

Item	Description
Project	The project name.
Cluster	The cluster name.
Service	The service name.
Server Role	The server role name.
APP	The application of the server role.

Item	Description
Name	The resource name.
Type	The resource type.
Status	The resource application status.
Parameters	The resource parameters.
Result	The resource application result.
Res	The resource ID.
Reprocess Status	The status of the interaction with Business Foundation System during the VIP resource application.
Reprocess Msg	The error message of the interaction with Business Foundation System during the VIP resource application.
Reprocess Result	The result of the interaction with Business Foundation System during the VIP resource application.
Refer Version List	The version that uses the resource.
Error Msg	The exception message.

1.5.10.10 Statuses of project components

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

Error State Component Table

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

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

Item	Description
Server Role Status	The current status of the server role.
Machine Status	The current status of the machine.

Server Role Alert Information

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

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

Machine Alert Information

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

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

Service Inspector Information

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

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

1.5.10.11 Relationship of service dependency

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

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

1.5.10.12 Check report of network topology

This report checks if network devices and machines have wirecheck alerts.

Check Report of Network Topology

Checks if network devices have wirecheck alerts.

Item	Description
Cluster	The cluster name.
Network Instance	The name of the network device.

Item	Description
Level	The alert level.
Description	The description about the alert information.

Check Report of Server Topology

Checks if servers (machines) have wirecheck alerts.

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

1.5.10.13 Clone report of machines

This report displays the clone progress and status of machines.

Clone Progress of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Status	The running status of the machine.
Clone Progress	The progress of the current clone process.

Clone Status of Machines

Item	Description
Project	The project name.
Cluster	The cluster name.
Machine Name	The machine name.
Machine Action	The action performed by the machine, such as the clone action.
Machine Action Status	The status of the action performed by the machine.
Machine Status	The running status of the machine.

Item	Description
Level	Whether the clone action performed by the machine is normal.
Clone Status	The current status of the clone action performed by the machine.

1.5.10.14 Auto healing/install approval pending report

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

1.5.10.15 Machine power on or off statuses of clusters

After a cluster starts or shuts down machines, you can view the related information in this report.

Cluster Running Statuses

If a cluster is starting or shutting down machines, the corresponding data is available in this list. No data indicates that no cluster has machines shut down.

Item	Description
Project	The project name.
Cluster	The cluster name.
Action Name	The startup or shutdown action that is being performed by the cluster.
Action Status	The status of the action.

Server Role Power On or Off Statuses

Displays the power on or off statuses of server roles in the cluster selected in the Cluster Running Statuses section.

Select a row in the Cluster Running Statuses section to display the information of the corresponding cluster in the list.

Item	Description
Cluster	The cluster name.
Server Role	The server role name.
Action Name	The startup or shutdown action that is being performed by the server role.
Action Status	The status of the action.

Statuses on Machines

Displays the running status of the selected server role on machines.

Select a row in the Server Role Power On or Off Statuses section to display the information of the corresponding server role in the list.

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

Machine Statuses

Displays the running statuses of machines in the selected cluster.

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

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

1.6 Network operations

1.6.1 Apsara Network Intelligence

1.6.1.1 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.

1.6.1.2 Log on to the Apsara Network Intelligence console

This topic describes how to log on to the Apsara Network Intelligence console.

Procedure

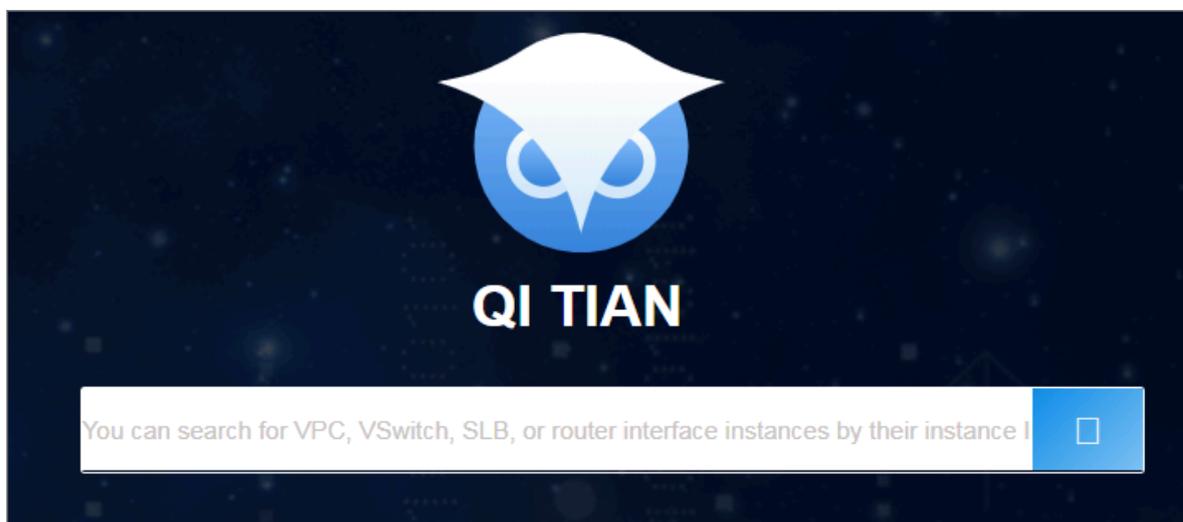
In the left-side navigation pane, click **Products**. On the right side of the page, click **Apsara Network Intelligence**.

1.6.1.3 Query information

You can enter an instance ID to query VPC, VRouter, and VSwitch details of the instance.

Procedure

1. [Log on to the Apsara Network Intelligence console.](#)



2. Enter the ID of a VPC or SLB instance to query instance details.

- **Enter a VPC instance ID to query VPC, VRouter, and VSwitch details.**

- **Instance details**

VPC Resources / VPC Details

Basic Information Subresource Information

Configuration Information

VPC ID	RegionNo	Status	Attached CENID	TunnelID
vpc-qbc44n...	cn-qingdao-em6-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	muyat_vpc	None	Yes
Enable ClassicLink	CIDR Block	User CIDR	Actions	
No	172.16.0.0/16	Details	Details	

- **Information about VRouters, routing tables, router interfaces, and VSwitches**

- **Enter the ID of an SLB instance to query instance details.**

- **Information about instance configurations, VIPs, specifications, and users**

VPC Resources / SLB Instance Details

Instance Information Listener Information

Configuration Information

LB ID	Cluster	EIP Type	Gateway Type	SLB Mode	status	
lb-qBk4...	cn-qingdao-em6-d01	intranet	classic	front	active	
VIPs	Proxies	Created At	Modified At	After WAF/Anti-DDoS Protection	Actions	
No data						
Cleaning Threshold	Black Hole Threshold					
None	None					
VIP(EIP) Information						
VIP(EIP)	Status	Tunnel ID	Service Unit Name	Primary IDC/VIS Name	Secondary IDC/VIS Name	
No data						
Specifications Information						
VIP MAX CONN LIMIT	VIP OUT bits/s	VIP IN bits/s	VIP QPS	VIP CPS	Specifications	Instance Type
No data						
User Information						
User ID	No data					

- **Listener information**

Click Show in the Back-end Server/Health Check column to view back-end server details.

VPC Resources / SLB Instance Details

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
lb-qBk4...	tcp	80	No	No	None	None	running	Show	2019-05-16 03:14:45	2019-05-16 03:14:56
lb-qBk4...	tcp	22	No	No	None	None	running	Show	2019-05-16 03:14:36	2019-05-16 03:14:56

1.6.1.4 Manage cloud service instances

You can create a cloud service in a region or query the instance information of a region.

Procedure

1. *Log on to the Apsara Network Intelligence console.*
2. **From the Products menu, choose Virtual Private Cloud > VPC Instance Type Management.**
3. **Select the region from the Select Region drop-down list for which you want to create a cloud service instance. All cloud service instances in the specified region are displayed.**
4. **Click Add to add a cloud service type.**

1.6.1.5 Tunnel VIP

1.6.1.5.1 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.

7. Click Create. On the Create SLB instance tab, select a primary data center or use the default data center.

8. Click OK. On the Create Listener tab, click Add to configure a UDP or TCP listener. Then, click Submit.

9. On the Publish Online tab, click Yes and click OK.

Result

The cloud services for which you have applied for VIPs can forward traffic through the created Layer-4 listener.

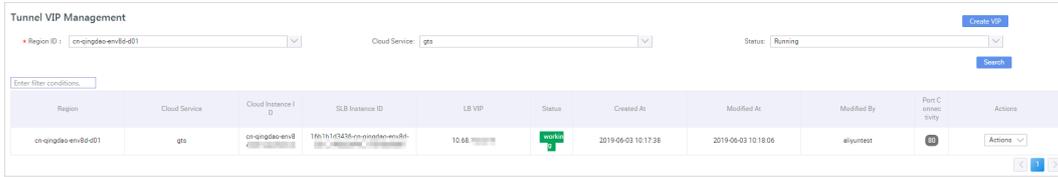
1.6.1.5.2 Query the tunnel VIP of a cloud service

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.

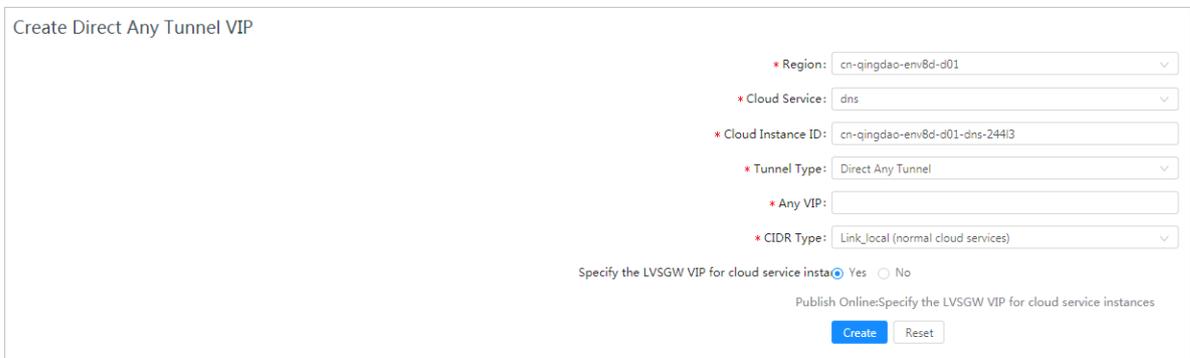


1.6.1.6 Create a Direct Any Tunnel VIP

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.

1.6.1.7 Leased line connection

1.6.1.7.1 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.

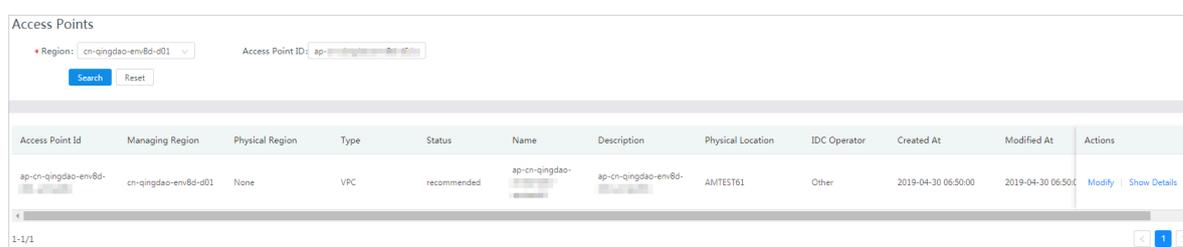
1.6.1.7.2 Manage an access point

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.



The screenshot shows the 'Access Points' management console. At the top, there is a search bar with a 'Region' dropdown set to 'cn-qingdao-em8-d01' and an 'Access Point ID' input field. Below the search bar are 'Search' and 'Reset' buttons. The main area contains a table with the following columns: Access Point Id, Managing Region, Physical Region, Type, Status, Name, Description, Physical Location, IDC Operator, Created At, Modified At, and Actions. One row is visible with the following data: Access Point Id: ap-cn-qingdao-em8-d01, Managing Region: cn-qingdao-em8-d01, Physical Region: None, Type: VPC, Status: recommended, Name: ap-cn-qingdao-em8-d01, Description: ap-cn-qingdao-em8-d01, Physical Location: AMTEST61, IDC Operator: Other, Created At: 2019-04-30 06:50:00, Modified At: 2019-04-30 06:50:00, and Actions: Modify | Show Details. At the bottom left, it shows '1-1/1' and at the bottom right, there are navigation icons.

Access Point Id	Managing Region	Physical Region	Type	Status	Name	Description	Physical Location	IDC Operator	Created At	Modified At	Actions
ap-cn-qingdao-em8-d01	cn-qingdao-em8-d01	None	VPC	recommended	ap-cn-qingdao-em8-d01	ap-cn-qingdao-em8-d01	AMTEST61	Other	2019-04-30 06:50:00	2019-04-30 06:50:00	Modify Show Details

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.

The screenshot shows a 'Modify Access Point' dialog box with the following fields and values:

- * 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:** (dropdown menu)

Buttons: **Modify** (blue), **Cancel** (white)

1.6.1.7.3 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

Region: Device ID:

Device ID	Region	Access Point ID	Device Status	Physical Location	Access Method	Device Name	Description	Created At	Modified At	Actions
CSW-VM-VPC-G1-...	cn-qingdao-em8-d01	ap-cn-qingdao-em8-d01-...	available	AMTEST61	vlanToVlanRouting	CSW-VM-VPC-G1-...	CSW-VM-VPC-G1-...	2019-04-29 22:50:32	2019-04-29 22:50:32	Modify Show Details

1-1/1

6. Click Show Details in the Actions column to view the details of the access device.

Modify access device information

Perform the following steps to modify the information about an access device:

1. Click **Modify** in the **Actions** column corresponding to a device that you want to modify.

2. Follow the on-screen prompts to modify the device information.

The screenshot shows a 'Modify Access Device' dialog box with the following fields and options:

- * Device ID:** CSW-VM-VPC-G: [blurred]
- * Region:** cn-qingdao-env8d-d01
- * Device Status:** Available Full Unavailable
- * Access Device Location:** AMTEST61
- * Specify whether to use XN:** Yes No
- * XNET Endpoint URL:** http://xnet.en[blurred]
- * XNET Device ID:** 1
- * Outer Source IP Encapsula:** 10.48[blurred]
- * Inner Source MAC Encapsu:** 00-00-5E-00-01-02
- Device Management IP Add:** 10.48.[blurred]
- Device Manufacturer:** Ruijie
- Device Model:** RG-S6220-[blurred]
- Device Name:** CSW-VM-VPC [blurred]
- Device Description:** CSW-VM-VPC-[blurred]

At the bottom of the dialog are two buttons: 'Modify' (highlighted in blue) and 'Cancel'.

3. Click Modify.

1.6.1.7.4 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.

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 > 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:

Description:

* BID:

* UID:

* Region:

The region ID is used for managing access devices (which is not necessarily the same as the attached region ID of the access device, but must be the same as the region ID of the access point).

* Access Point Type: VPC Access Point

- VPC -VPC access point, for leased lines that can access VPC networks

* Access Point ID:

Access Point ID

Device Name:

Device Port:

Bandwidth: Mbps

The inbound interface bandwidth of the leased line. Unit: Mbit/s. Value range: [2-10000].

* Port Type:

You can leave it empty if the value is unknown.

Redundant Leased Lines:

- When establishing the second leased line, you can specify it as a redundant one and upload its ID. If you do so, Alibaba Cloud allocates a separate access device for higher availability.
- The leased line that you specify must exist and be in Allocated, Confirmed, or Enabled status.

Create
Cancel

When the leased line state is Confirmed, the line is created.

5. On the Leased Lines page, find the created leased line and choose Actions > Enable.

If the allocation process for a leased line persists for several minutes after you click Enable, choose Products > Network Controller > Business Foundation System Flow. On the page that appears, set Instance ID to the leased line ID, set Step Status to All, and click Search. Check the flow status in the search results. A flow in red indicates that the corresponding step has failed. Click Resend to restart the task, and then requery the flow status.

If the flow fails, run the `vpcregiondb -e "select * from xnet_publish_task order by id desc limit 5"` command on the ECS availability group (AG). If an error is returned, you can check the xnet service logs to troubleshoot the issue based on the returned error.

1.6.1.7.5 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.

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 > VBRs.

4. Click Create VBR.

Create VBR
✕

* BID:

* UID:

* Region: ▼
 The ID of the region to which the instance belongs.

* Leased Line ID:

* VLAN ID:
 The VLAN of the VBR leased line interface.

- VLAN : [1, 2999]
- Only the leased line owner can specify or modify VLAN.

* Local Gateway IP Address:

- The local IP address of the leased line interface.
- It is required when the interface status is not waiting.
- Only the VBR owner can specify or modify the local IP address.

* Peer Gateway IP 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:

- The subnet mask for the connection between the local IP addresses 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:
 The leased line name. It can be 2 to 128 characters in length and cannot start with http:// or https://.

Description:
 The leased line description. It can be 2 to 128 characters in length and cannot start with http:// or https://.

ownerBid:

ownerAliUid:

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 Information ————— ② Peer Information ————— ③ Results

Select Router Type: Single Double

Name:

Description:

* Bid:

* Uid:

* Region: ▼

* Router Type: VRouter VBR

Zone:

* Router ID:

* Role: InitiatingSide AcceptingSide

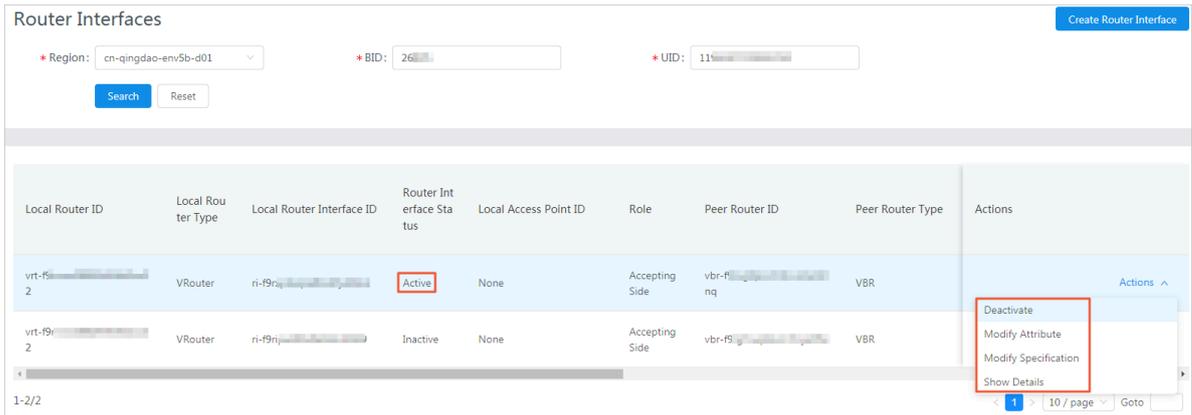
* Specifications: ▼

Health Check Source IP:

Health Check Destination:

Skip Inventory Check: Yes No

When the router interface state is Active, the interface is created.



1.6.1.7.7 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.

Procedure

1. Perform the following steps to add routes on a VBR destined for a VPC and an IDC:

- a) *Log on to the Apsara Network Intelligence console.*
- b) **From the Products menu, choose Express Connect > Network Environment Management.**
- c) **Choose Function Modules > Routing Tables.**
- d) **Set search conditions such as Region, BID, UID, Router Type, Routing Table ID, and Router ID, and click Search to query routing tables.**
- e) **Click Add Route Entry in the Actions column corresponding to a routing table.**
- f) **Specify a route entry destined for the CIDR block of a destination VPC, and click Create.**

The parameters are described as follows:

- **Destination CIDR Block:** the destination CIDR block.
- **Next Hop Type:** the next hop type.

- **Next Hop Instance ID:** the ID of the next hop instance for the specified next hop type.

Figure 1-11: Add a route destined for a destination VPC

Add Routing Entry
✕

* BID:

* UID:

* Routing Table ID:

Modify the routing table ID to which the routing entry belongs.

* Destination CIDR Block:

The network mask, such as 255.255.255.0/24.

* ECMP: Yes No

* Next Hop Type:

- The next hop type. Valid values: Instance, Tunnel, HaVip, RouterInterface.
- Set the value to RouterInterface for ECMP.

* Next Hop ID:

The next hop interface ID for the route entry.

Create
Cancel

g) Repeat the preceding steps to add a route destined for a target IDC.



Note:

You can navigate to the VBRs page and locate the VLAN Interface ID area to obtain next hop router interface information.

2. Add a route destined for the router interface of a VBR in the VPC.
3. On the gateway of the on-premises IDC, configure a route destined for the VPC.

1.6.1.8 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.

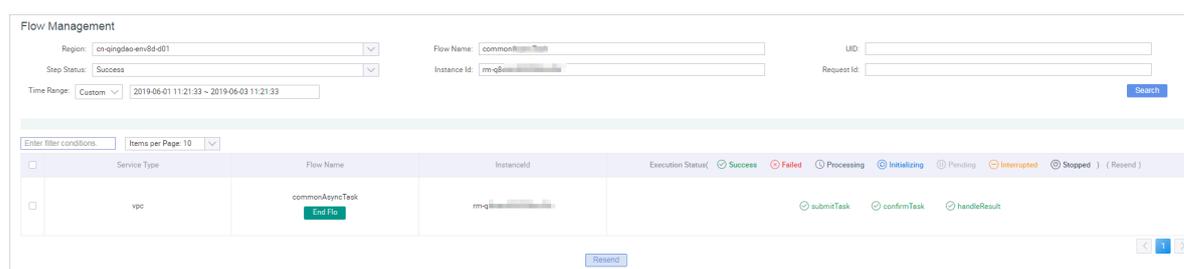
Procedure

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.

Figure 1-12: Flow Management page



1.6.1.9 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.

Procedure

1. *Log on to the Apsara Network Intelligence console.*
2. From the Products menu, choose Cloud Service Management > Cloud Service Reverse Access.
3. On the page that appears, enter AccessKey ID and AccessKey Secret and click OK. The Cloud Service Reverse Access page appears.
4. Click Create Cloud Service Reverse Access.
5. On the Allocate App ID tab, set Region, Name, and Description.
6. Click Continue. The following information is automatically created and displayed on the Create Address Pool tab: the application IDs of cloud services that allow reverse access and the address pools that are used for reverse access to the cloud services.

7. **Click Continue. On the Add Server Address tab, configure an ECS instance to be used for reverse access.**
 - **VPC ID:** specifies the ID of a VPC, an ECS instance, or a single-tunnel cloud service instance.
 - **Server IP:** specifies the IP address of the ECS instance to be used for reverse access.
8. **Click Continue. On the Create Mapping IP tab, configure VSwitch ID and Mapping IP of the ECS instance in the destination VPC.**
9. **Click Continue. On the Complete Authorization tab, configure VPC ID, ECS Instance IP, and Instance Port for reverse access.**

The value of Instance Port must be an integer value. You can specify multiple instance ports separated by commas (.). Example: 10,20,30. You can configure up to 10 instance ports.

2 Operations of basic cloud products

2.1 Elastic Compute Service (ECS)

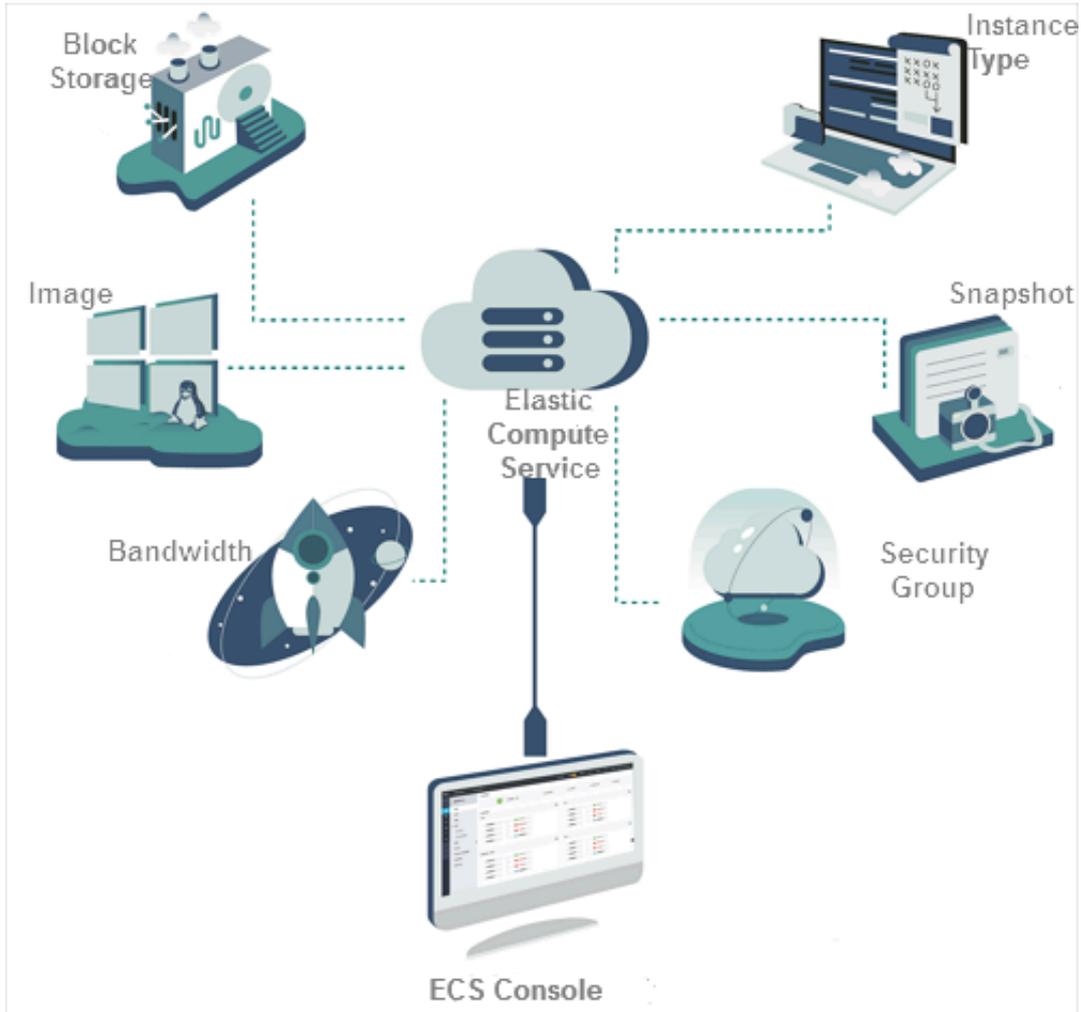
2.1.1 ECS overview

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 [Figure 2-1: ECS instance](#).

Figure 2-1: ECS instance



2.1.2 Log on to Apsara Stack Operations

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

Prerequisites

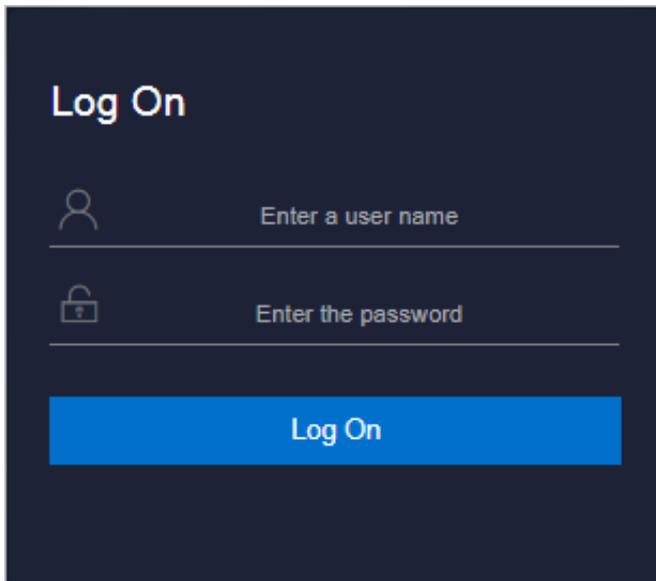
- ASO access address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

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.

Figure 2-2: Log on to ASO



Note:

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

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

2.1.3 ECS operations and maintenance

2.1.3.1 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.

2.1.3.2 VM

2.1.3.2.1 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.

2.1.3.2.2 Search for VMs

In Apsara Stack Operations, you can view the list of existing VMs and their information.

Procedure

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

2. In the left-side navigation pane, choose **Products > ECS**.

The ECS Operations and Maintenance Platform is displayed.

3. On the VMs tab, set search conditions, and click **View** to search for the specified VMs.

Region is a required search condition.

4. In the VM list, click a VM ID, and the VM Details page appears on the right side of the page.

2.1.3.2.3 Start a VM

You can start a VM through Apsara Stack Operations. The operations are similar to those on a real server.

Prerequisites

A VM must be in the `stopped` state to be started.

Procedure

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

2. In the left-side navigation pane, choose **Products > ECS**.

3. On the ECS Operations and Maintenance Platform page that appears, set search conditions, and click View to search for the specified VMs.
4. In the VM list, select the VM you want to start, and click Start above the list.
5. In the Start VM dialog box that appears, you can set Start to Normal or Repair.



Note:

If you need to reset the network settings for the VM, set Start to Repair. Otherwise, set Start to Normal.

6. Set Operation Reason and click OK.

2.1.3.2.4 Stop a VM

You can stop a VM through Apsara Stack Operations. The operations are similar to those on a real server.

Prerequisites

- A VM must be in the running state to be stopped.
- 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 > ECS.
3. On the ECS Operations and Maintenance Platform page that appears, set search conditions, and click View to search for the specified VMs.
4. In the VM list, select the VM you want to stop, and click Stop above the list.
5. In the Stop VM dialog box that appears, you can set Shutdown Policy to 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.

6. Set Operation Reason and click OK.

2.1.3.2.5 Restart a VM

You can restart a VM through Apsara Stack Operations. The operations are similar to those on a real server.

Prerequisites

- A VM must be in the running state to be restarted.
- 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 > ECS**.
3. On the ECS Operations and Maintenance Platform page that appears, set search conditions, and click **View** to search for the specified VMs.
4. In the VM list, select the VM that you want to restart, and click **Reboot** above the list.
5. In the Reboot VM dialog box that appears, set **Start and Shutdown Policy** as needed.



Note:

- For the **Start** parameter, you can select **Normal** or **Repair**.
 - For the **Shutdown Policy**, you can select **Non-force Shutdown** or **Force Shutdown**.
6. Set **Operation Reason** and click **OK**.

2.1.3.2.6 Cold migration

In Apsara Stack Operations, you can migrate a VM while it is offline to implement failover.

Prerequisites

- Failover must be performed offline. Make sure that the VM is in the stopped state before you start cold migration.
- Failover can only be performed within the same zone. Cross-zone failover is not allowed.

Context

In the event that a VM or an NC fails, you must fail over the VM by shutting the VM down and migrating it to a new NC. This is an example of failover.

Procedure

1. *Log on to Apsara Stack Operations.*
2. **In the left-side navigation pane, choose Products > ECS.**
3. **On the ECS Operations and Maintenance Platform page that appears, set search conditions, and click View to search for the specified VMs.**
4. **In the VM list, select the VM that needs to be migrated, and click Stop and Migrate above the list.**
5. **In the dialog box that appears, select NC Switching, and then set Switching Policy, Start, and Recovery.**
6. **Set Operation Reason and click OK.**

2.1.3.2.7 Reset a disk

You can reset disks to restore them back to their initial state when necessary.

Prerequisites

- **When you reset a disk, any applications that are installed to the created VM will be lost. Before you perform a reset operation, ensure that your data is backed up.**
- **To reset a disk, the VM to which it belongs must be 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 > ECS.**
3. **On the ECS Operations and Maintenance Platform page that appears, set search conditions, and click View to search for the specified VMs.**
4. **In the VM list, select the VM that contains the disk to be reset.**
5. **Choose More > Reset Disk.**
6. **In the dialog box that appears, select the disk to be reset, set Operation Reason, and click OK.**

2.1.3.3 Disks

2.1.3.3.1 Overview

In an ECS instance, cloud disks can be considered as physical disks. You can mount, detach, and create snapshots for disks.

2.1.3.3.2 Search for disks

You can view a list of existing disks and their information in Apsara Stack Operations.

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page that appears, click the **Disks** tab.
4. On the Disks tab, set search conditions, and click **View** to search for the specified disks.
Region is a required search condition.

2.1.3.3.3 Search for snapshots

You can view the list of existing snapshots and their information in Apsara Stack Operations.

Procedure

1. *Log on to Apsara Stack Operations.*
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page that appears, click the **Disks** tab.
4. On the Disks tab, set search conditions, and click **View** to search for the specified disks.
5. Click the  icon next to the disk whose snapshots you want to view, and click **View Snapshot**.
The information of all snapshots on the disk is displayed.

2.1.3.3.4 Mount a disk

After a disk is created, you must mount the disk. Only independent cloud disks can be mounted to ECS instances.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page that appears, click the **Disks** tab.
4. On the Disks tab, set search conditions, and click **View** to search for the specified disks.
5. Click the  icon next to the disk that you want to mount, and click **Mount**.
6. In the dialog box that appears, set **VM ID** and **Operation Reason**, and then click **OK**.

2.1.3.3.5 Detach a disk

Only Apsara Stack Operations can be used to detach data disks. System disks and local disks cannot be detached.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page that appears, click the **Disks** tab.
4. On the Disks tab, set search conditions, and click **View** to search for the specified disks.
5. Click the  icon next to the disk that you want to detach, and click **Detach**.
6. In the dialog box that appears, set **Operation Reason** and click **OK**.

2.1.3.3.6 Create a snapshot

You can manually create a snapshot of a disk as needed in Apsara Stack Operations.

Procedure

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

2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance Platform** page, click the **Disks** tab.
4. On the **Disks** tab, set search conditions, and click **View** to search for the specified disks.
5. Click the  icon next to a disk, and select **Take Snapshot**.
6. In the dialog box that appears, set **Snapshot Name**, **Snapshot Description**, and **Operation Reason**. Click **OK**.

2.1.3.4 Snapshots

2.1.3.4.1 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.

2.1.3.4.2 Search for snapshots

You can view the list of existing snapshots and their information in **Apsara Stack Operations**.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the **ECS Operations and Maintenance Platform** page, click the **Snapshots** tab.
4. On the **Snapshots** tab, set search conditions, and click **View** to search for the specified snapshots.

Region and **AliUid** are required search conditions.

2.1.3.4.3 Delete a snapshot

In Apsara Stack Operations, you can delete snapshots that are no longer used.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page, click the **Snapshots** tab.
4. On the Snapshot tab, set search conditions, and click **View** to search for the specified snapshots.
5. Click the  icon next to a snapshot, and click **Delete**.
6. In the dialog box that appears, set **Operation Reason** and click **OK**.

2.1.3.4.4 Create an image

You can create a custom image by using a snapshot. The operating system and environment variables of the snapshot are contained in the image.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page that appears, click the **Snapshots** tab.
4. On the Snapshots tab, set search conditions, and click **View** to search for the specified snapshots.
5. Click the  icon next to a snapshot, and click **Create Image**.
6. In the dialog box that appears, set **Image Name**, **Image Version**, **Image Description**, and other parameters. Click **OK**.

2.1.3.5 Images

2.1.3.5.1 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.

2.1.3.5.2 Search for images

You can view the list of existing images and their information in Apsara Stack Operations.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page, click the **Images** tab.
4. On the **Images** tab, set search conditions and click **View**.

Region is a required search condition.

2.1.3.6 Security groups

2.1.3.6.1 Overview

A security group is a virtual firewall that provides Stateful Packet Inspection (SPI). Security groups provide virtual firewall-like functionality and are used for network access control for one or more ECS instances. They are important means of network security isolation and are used to divide security domains on the cloud.

Security group rules can permit the inbound and outbound traffic of the ECS instances associated with the security group. You can authorize or cancel security group rules at any time. Changes to security group rules are automatically applied to ECS instances that are members of the security group.

When you configure security group rules, ensure that the rules are concise and easy to manage. If you associate an instance with multiple security groups, hundreds of rules may apply to the instance, which may cause connection errors when you access the instance.

2.1.3.6.2 Search for security groups

You can view the list of current security groups and their information in Apsara Stack Operations.

Context

You can modify security group rules to permit or deny public or internal network traffic to and from the ECS instances associated with the security groups. You can

add or delete security group rules at any time. Changes to security group rules are automatically applied to ECS instances that are members of the security group.



Note:

- If two security groups have the same security group rules but different access rules, deny rules will take precedence over allow rules.
- No rule in a security group can permit outbound access from an ECS instance while forbidding inbound access to the ECS instance.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page, click the **Security Groups** tab.
4. On the Security Groups tab, set search conditions and click **View**.
Region is a required search condition.
5. Click **Add Rule** to add a rule to this security group.

2.1.3.6.3 Add security group rules

You can add rules to security groups for your service operations.

Procedure

1. [Log on to Apsara Stack Operations](#).
2. In the left-side navigation pane, choose **Products > ECS**.
3. On the ECS Operations and Maintenance Platform page, click the **Security Groups** tab.
4. On the Security Groups tab, set search conditions and click **View**.
5. Click the  icon next to a security group, and click **Add Rule**.

6. In the dialog box that appears, configure the parameters.

Table 2-1: Security group rule parameters describes the parameter configurations.

Table 2-1: Security group rule parameters

Parameter	Description
Protocol	<ul style="list-style-type: none"> • TCP • UDP • ICMP • GRE • ALL: All protocols are supported.
Rule Priority (1 to 100)	The smaller the value, the higher the priority.
Network Type	<ul style="list-style-type: none"> • Internet: the public network • Intranet: the internal network
Authorization Policy	<ul style="list-style-type: none"> • Accept: accepts the packet on access. • Drop: abandons the packet on access. • Reject: denies the packet on access.
Port Number Range	1 to 65535, such as 1/200, 80/80, or -1/-1.
Access Direction	<ul style="list-style-type: none"> • Ingress: Inbound traffic is allowed for the security group. • Egress: Outbound traffic is allowed for the security group.
IP Address Range	If the authorization type is IP address segment access, the authorized user needs to 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 are supported.
Security Group ID	Enter the ID of the associated security group.
Operation Reason	Optional. You can enter related operation reasons.

7. When the parameter configuration is completed, click OK.

2.1.4 VM hot migration

2.1.4.1 Overview

Hot migration is the process of migrating a running VM from one host to another.

During migration, the VM runs normally and its services are not aware that any

migration task is occurring. However, these services can detect a very short interruption between 100 and 1,000 ms.

Scenarios

During system operations and maintenance, hot migration is typically used for the following scenarios:

- **Active O&M:** The host is faulty and must be repaired, but the fault does not affect the operation of the system. You can use hot migration to migrate the VM to another host and repair the faulty host in offline mode.
- **Server load balancing:** When a host is experiencing a high load, you can migrate some of its VMs to other idle hosts to reduce resource consumption on the source host.
- **Other scenarios** where a VM must be migrated without affecting its business operations.

2.1.4.2 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.

2.1.4.3 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_check_image] [no_check_load] [downtime]==
Usage: == houyi_api.sh <function_name> [--help|-h] [name=value]
```

Table 2-2: Parameter description

Parameter	Function	Impact	Value
<code>vm_name</code>	The name of the VM to be migrated.	N/A	N/A
<code>nc_id</code>	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
<code>rate</code>	The amount of host bandwidth to be allocated for migration tasks.	The migration will use the bandwidth resources of the hosts.	<ul style="list-style-type: none"> 10 GB network: 80 MB 1 GB network: 40 MB
<code>downtime</code>	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

Parameter	Function	Impact	Value
no_check_image	Forcibly migrates the images that are not supported.	Performing this operation may violate the SLA.	false
no_check_load	Forcibly migrates images even when the load threshold requirements are not met.	Downtime cannot be controlled when this parameter is set to false.	false

Cancel hot migration

Run the following command to cancel a hot migration task:

```
go2hyapi cancel_live_migrate_vm == Usage: == houyi_api.sh <
function_name> [--help|-h] [name=value] == Functions usage: == |-
cancel_live_migrate_vm <region_id> <vm_name>
```

Table 2-3: Parameter description

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
region_id	The ID of the region where the target VM is located.	N/A	N/A

2.1.4.4 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_change_migration_status <vm_name> <
region_id> <where>
```

Table 2-4: Parameter description

Parameter	Function	Impact	Value
vm_name	The name of the VM to be migrated.	N/A	N/A
region_id	The ID of the region where the target VM is located.	N/A	N/A
where	The ID of the NC where the VM is located.	N/A	N/A

2.1.4.5 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
 - nc_id
- What preparations should I make before performing a hot migration operation?
 - Confirm that the VM is in the running state.
 - Confirm the destination of the VM migration.

- **Can hot migration be canceled? How can I cancel hot migration?**

Yes. If the API request is successful and the migration has not completed, run the `go2hyapi cancel_live_migrate_vm vm_name=[vm_name] region_id=[region_id]` command to cancel the hot migration. If the VM has completed its migration to the destination NC, it is too late to cancel the hot migration.

You can get the value of `region_id` by running the `go2which [vm_name]` command to view `region_info`.

- **The VM is still in the migrating state after the hot migration has completed, and the `cancel_live_migrate_vm` command is not working. What should I do?**

You can run the `virsh query-migrate [domid]` command on the source NC of the VM to check whether the VM is still being migrated. If the VM is still being migrated, a piece of JSON information will be returned. If the VM has finished migration, run the following command on the AG to modify the state of the VM:

```
go2hyapi manually_change_migration_status vm_name=[vm_name] where=[nc_id for the VM] region_id=[region_id]
```

`domid` is the name of the VM instance. You can run the `virsh list|grep vm_name` command to view it.

- **How can I confirm whether the VM is migrated successfully?**

On the destination NC of the VM, run the `sudo virsh list|grep [vm_name]` command. If the VM instance exists and is not in the running state, the migration is successful.

- When an exception occurs during hot migration, which logs should I refer to?

- View the Libvirt bottom layer migration log on the NC.

Run the `/var/log/libvirt/libvirt.log` command to view information about the migration process, such as vport offline, detach, delete, and relay route.

- Run the following command to view the API management log of Server Controller on the AG:

```
/var/log/houyi/pync/houyipync.log
```

- View the Qemu log.
- Run the following command to view the regionmaster log on the VM:

```
regionmaster/logs/regionmaster/error.log
```

- A VM fails to start after hot migration. Is the VM still in the pending state?

If error vport update nc conf by vpc master fails dest_nc_id:xxx is returned, it indicates that a VPC fault has occurred and the underlying task is interrupted.

- During hot migration, the API returns the following error message: distributed lock fail. What are the possible causes of this issue?

The API has been called too many times within a short period of time. Wait several minutes and then try again.

- What are some common scenarios where migration fails? How can I resolve these issues?

Table 2-5: 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 <code>no_check_load=true</code> to skip this inspection.
The VM fails to pass image inspection.	It is not an Alibaba Cloud-specified image.	You can run <code>no_check_image=true</code> to skip this inspection. Be aware of the risks involved.

2.1.5 Hot migration of disks

2.1.5.1 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.

2.1.5.2 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* .

2.1.5.3 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 ($allTaskIds) and status=2 and src_recycled=0 and DATE(gmt_finish) < DATE_ADD(CURDATE(), INTERVAL -1 DAY)'` command.
3. Perform the following operations for each set of `<task_id,src_pangu_path,dst_pangu_path>`:
 - a) Run the `/apsara/deploy/bsutil rlm --dir=$dst_pangu_path|grep 'not-loaded'|wc -l` command on the host that runs the `bstools` role in the storage cluster. If the command output is not 0, proceed to the next step.
 - b) Run the `/apsara/deploy/bsutil delete-image --dir=$src_pangu_path` command on the host that runs the `bstools` role in the storage cluster.
 - c) Run the `/apsara/river/river_admin migrate recycle $task_id` command on the host that runs the `river` role in the storage cluster.

2.1.6 Upgrade solution

2.1.6.1 Overview

For both hot and cold migration of GPU and FPGA clusters, you must understand the limitations that apply to cluster upgrades.

2.1.6.2 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.**

2.1.6.3 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.**

2.1.7 Disk maintenance of an instance

2.1.7.1 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.

2.1.7.2 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.

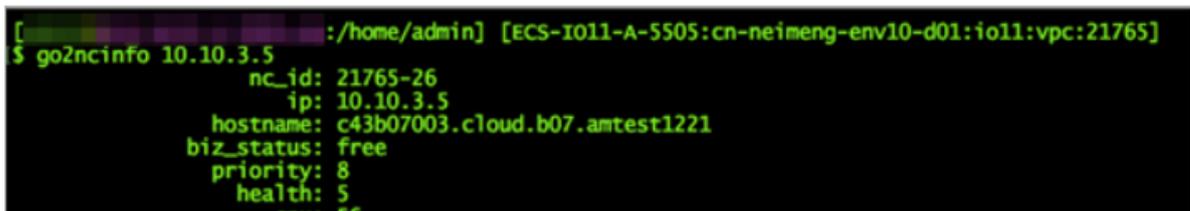
Run the following command to obtain the NC ID based on the NC IP address:

```
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
- Run the go2ncinfo 10.10.3.5 command to obtain the NC ID.
- NC ID: 21765-26



```
[root@ecs-1011-a-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ go2ncinfo 10.10.3.5
nc_id: 21765-26
ip: 10.10.3.5
hostname: c43b07003.cloud.b07.amtest1221
biz_status: free
priority: 8
health: 5
```

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-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ go2hyapi query_vm_list format=json region_id=cn-neimeng-env10-d01 nc_id=21765-26 nc_storage_device_id=/apsarapangu/disk1
[ERROR] [2018-05-10 16:41:36] The function 'query_vm_list' doesn't exist!
Usage:
houyi_api.sh <function_name> [name=value]
available functions:
```

If an error is reported when the API is used, you must run the following command instead. The local_disk_ops.py script is in the /home/admin directory in this environment.

```
/home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk --logfile=/tmp/tmp.log nc_id=21765-26 storage_device_id=/apsarapangu/disk1
```

```
[admin@ ~ :/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ ls | grep local
local_disk_ops.py
[admin@ ~ :/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ /home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk --logfile=/tmp/tmp.log nc_id=21765-26 storage_device_id=/apsarapangu/disk1
[{'vm_name': 'i-5wf05ykw7mic5aq65dv2', 'status': 'running'}]
```

You can see that only the i-5wf05ykw7mic5aq65dv2 instance runs on this disk and is in the running state.

3. Shut down the VMs on the AG by using Server Controller.

- a) If the VMs are in the running state, you need to shut them down first.

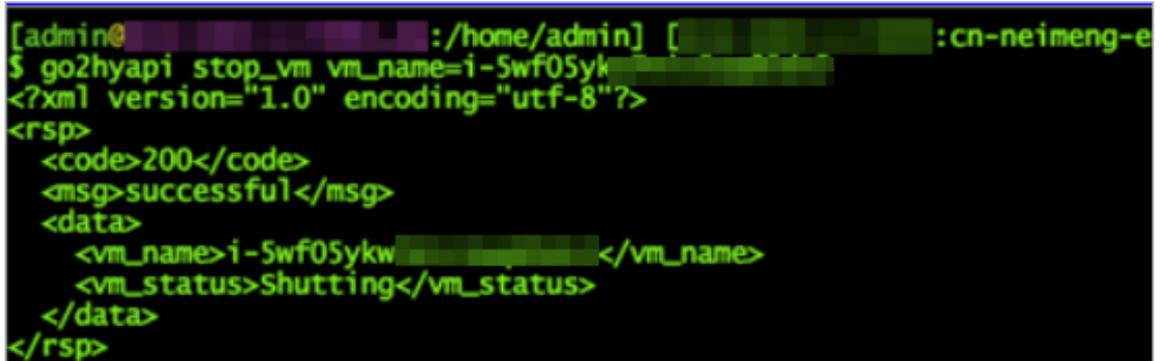
Run the following command:

```
go2hyapi stop_vm vm_name={vm_name}
```

{vm_name} is the ID of the running VM obtained in the preceding step.

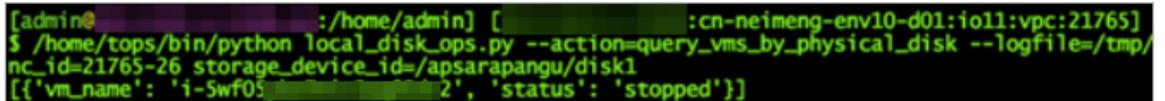
Example:

```
go2hyapi stop_vm vm_name=i-5wf05ykw7mic5aq65dv2
```



```
[admin@ :/home/admin] [ :cn-neimeng-e
$ go2hyapi stop_vm vm_name=i-5wf05ykw
<?xml version="1.0" encoding="utf-8"?>
<rsp>
  <code>200</code>
  <msg>successful</msg>
  <data>
    <vm_name>i-5wf05ykw</vm_name>
    <vm_status>Shutting</vm_status>
  </data>
</rsp>
```

Wait until the VM status changes to Stopped.



```
[admin@ :/home/admin] [ :cn-neimeng-env10-d01:io11:vpc:21765]
$ /home/tops/bin/python local_disk_ops.py --action=query_vms_by_physical_disk --logfile=/tmp/
nc_id=21765-26 storage_device_id=/apsarapangu/disk1
[{'vm_name': 'i-5wf05ykw7mic5aq65dv2', 'status': 'stopped'}]
```

- b) If the VM is in the pending or stopped state, you do not need to shut it down.
- c) If the VM is in another state, you must wait until its status changes to running, pending, or stopped. Alternatively, you can carry out an inspection.

4. Use Server Controller to check the local data disk associated with the physical disk.

Run the following command on the AG:

```
$/home/tops/bin/python local_disk_ops.py --action=query_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=/tmp/tmp.log nc_id=21765-26 storage_device_id=/apsarapangu/disk1
```

```
[admin@ ~ :/home/admin] [IP: 10.10.10.10] [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...', '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.

a) Check the device file name of the damaged disk on the NC.

Run the following command on the NC:

```
df -h
```

Example:

The device file name corresponding to /apsarapangu/disk1 is /dev/sdb1.

b) Check the serial numbers (SN) of the NC and the hard disk.

A. In the Apsara Infrastructure Management Framework console, check the SN of the NC in the corresponding cluster operation and maintenance center.

The SN of the NC is used to locate the machine if the disk is replaced on site.

Example: CVXKB7CD00J

B. Check the SN of the hard disk.

Run the following command:

```
smartctl -a {device_file_name} | grep 'Serial Number'
```

{device_file_name} is the device file name obtained earlier.

Example:

```
smartctl -a /dev/sdb1 | grep 'Serial Number'
```



```
[root@cloud-ops:~]# cd /proc/scsi
[root@cloud-ops:/proc/scsi]# smartctl -a /dev/sdb1 | grep 'Serial Number'
Serial Number:      K1K3EPKD
```

The SN of /dev/sdb1: K1K3EPKD

c) Remove the original disk.

The on-site engineer 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	CVXKB7CD00T	N41S1-6T. 22
------------	-----	----	-------------	--------------

The N41S1-6T model supports hot swapping and uses the M.2 card as its system disk. The 12 hard disks can be seen on the front panel.

The disk order is as follows:

- /dev/sdb : 1 /dev/sde : 4 ...
- /dev/sdc : 2 /dev/sdf : 5 ...
- /dev/sdd : 3 /dev/sdg : 6 ...



You need to remove the /dev/sdb1 hard disk from slot 1. The SN of the hard disk should be consistent with the K1K3EPKD SN obtained earlier.



- d) Insert a new disk.
- e) Partition and mount the disk, and modify the label and the fstab file. The new disk must be mounted to the original mount point.

A. Check whether the hard disk is installed correctly.

Run the `fdisk -l` command to view the ID of the hard disk.

Example:

```

Disk /dev/sdb: 6001.2 GB, 6001175126016 bytes, 11721045168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 4096 bytes
I/O size (minimum/optimal): 4096 bytes / 4096 bytes
Disk label type: dos
Disk identifier: 0x00000000

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1             1  4294967295  2147483647+  ee    GPT
Partition 1 does not start on physical sector boundary.

```

You can see that the new hard disk is identified as sdb.

B. Partition the hard disk.

Run the `fdisk` command if the hard disk capacity is not greater than 2 TB.

```
fdisk /dev/sdb
```

Run the `parted` command if the hard disk capacity is greater than 2 TB.

```
parted /dev/sdb
```

The `parted` command is used to partition the 5.5 TB hard disk.

```
mklabel gpt
```

Use the GPT to form a 5.5 TB partition.

```

(parted) mklabel gpt
warning: The existing disk label on /dev/sdb will be destroyed and all data on this disk will be
lost. Do you want to continue?
Yes/No? Yes

```

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@ ~]# lsblk | grep sdb
sdb      8:16    0    5.5T    0 disk
└─sdb1   8:17    0    5.5T    0 part

```

C. Format the partition.

```
mkfs -t {filesystem_type} {device_name}
```

{filesystem_type} is the type of the file system to be formatted. {device_name} is the name of the partition to be formatted.

Example:

```
mkfs -t ext4 /dev/sdb1
```

```
[root@ ~ :/root]
#mkfs.ext4 /dev/sdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
183144448 inodes, 1465130240 blocks
73256512 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=3613392896
44713 block groups
32768 blocks per group, 32768 fragments per group
4096 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000, 214990848, 512000000, 550731776, 644972544

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

```
[root@ ~ :/root]
#lsblk -T
NAME        FSTYPE LABEL        UUID                                 MOUNTPOINT
sda
├─sda1     ext4  /boot        1fd12aa3-8f54-4bb0-a1d3-a29595f391b8 /boot
├─sda2     ext4  /            3ac491f4-c2a4-4372-a4c3-3b3605b8a6da /
├─sda3     swap  SWAP         57955bd2-1038-4f7e-8e85-f3b16d95794d
├─sda4
├─sda5     ext4  /apsarapangu 0e287a91-1e95-47a9-a815-c9a6b80d821e /apsarapangu
├─sda6     ext4  /ansara      67aec0b4-9bd0-4601-96ea-973d006c0979 /ansara
└─sdb
   └─sdb1   ext4  disk1       fd10be5a-efac-4cc7-8ecd-f1dd8df7824d
sdc
├─sdc1     ext4  disk2       a3b778c6-dc3e-40fe-89fe-6593d48db54e /apsarapangu/disk2
sdd
├─sdd1     ext4  disk3       b7c2c0c3-379d-41f2-9a09-dbc6add14093 /apsarapangu/disk3
sde
├─sde1     ext4  disk4       369a120f-4cd0-4249-b6cb-17c995a662cc /apsarapangu/disk4
sdf
```

D. Mount the hard disk to 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.

```
[root@ ~]# cat /etc/fstab | grep 'disk1'
LABEL=disk1 /apsarapangu/disk1 ext4 noatime,nodiratime,nobarrier 0 0
```

```
e2label /dev/sdb1 disk1
```

```
[root@ ~]# blkid
/dev/sdb1: LABEL="disk1" UUID="65ce9f79-ab6f-48ea-8e28-84a2bb3ff420" TYPE="ext4" PARTLABEL="primary" PARTUUID="6a0c5246-1002-4b5b-be24-c8d2ae20eff3"
```

F. 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.

```
[root@ :/apsarapangu/disk1]
#ls
3388  lost+found
```

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
```

```
[admin@ :/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$/home/tops/bin/python local_disk_ops.py --action=reset_local_disk_after_change_physical_disk --logfile=/tmp/tmp.log disk_id=1000-3388
OK
```

OK indicates 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
```

```
[admin@ :/home/admin] [ECS-I011-A-5505:cn-neimeng-env10-d01:io11:vpc:21765]
$ go2hyapi start_vm vm_name=i-5wf05ykw7mic5aq65dv2
<?xml version="1.0" encoding="utf-8"?>
<rsp>
  <code>200</code>
  <msg>successful</msg>
  <data>
    <vm_name>i-5wf05ykw7mic5aq65dv2</vm_name>
    <vm_status>Starting</vm_status>
  </data>
</rsp>
```

Result

You can log on to the VM through SSH, format the device corresponding to the new disk, and mount it to the mount point. Check the disk capacity and whether data read/write operations are successful.

2.1.7.3 Additional instructions

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](#).

2.1.8 Handle routine alarms

2.1.8.1 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.

Table 2-6: Description of metric types

Metric type	Function	Solution
Basic metric/ service availability metric	Monitors the basic performance of the host and the availability of the services on the host. This kind of metrics includes CPU, memory, and handle count.	<p>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.</p> <p>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.</p>
Connectivity metric	Checks the connectivity between each module and its related modules.	<ul style="list-style-type: none"> • 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.	<ul style="list-style-type: none"> • 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.

2.1.8.2 API proxy

This topic describes the metrics of API proxy.

Table 2-7: Metric description

Metric	Alert item	Description
check_apip roxy_dns	Database HA switchover occurs or not	Checks whether Server Controller database switchover occurs. If so, nginx will be reloaded automatically.
check_apip roxy_conn_new	check_apip roxy_conn_new	Checks the connectivity to the Server Controller database.
		Checks the connectivity to the API Server: <ul style="list-style-type: none"> · Checks whether the API Server is down. · Checks the network connectivity.
check_apip roxy_proc_new	check_apip roxy_proc_new	Checks the memory usage and CPU utilization for nginx and memcache processes.

2.1.8.3 API Server

The topic describes the metrics of the API Server.

Table 2-8: Metric description

Metric	Alert Item	Solution
check_API Server_proc_new	The process does not exist or is abnormal.	Checks the state of the Java process: whether the process exists, and the CPU utilization and memory usage

Metric	Alert Item	Solution
check_API Server_con n_new	Checks the connectivity between the API Server and Server Controller database.	Checks whether the corresponding component is down. If the corresponding component is down, fix the issue by taking necessary O&M measures. If the database is down, contact DBA to fix the issue. Checks whether the VIP is connected to the corresponding component. If not, contact the network engineer to fix it.
	Checks the connectivity between the API Server and TAIR.	
	Checks the connectivity between the API Server and RegionMaster.	
	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.	<ul style="list-style-type: none"> • If an exception occurs to the database, contact DBA to check whether the database is normal. • If the creation of an instance fails, locate the cause of the failure.

2.1.8.4 RegionMaster

This topic describes the metrics of RegionMaster.

Table 2-9: Metric description

Metric	Alert item	Description
check_regionmaster_p roc	The process does not exist or is abnormal.	Checks the state of the Java process: whether the process exists, and the CPU utilization and memory usage.
check_regionmaster_w ork	rms_connectivity	Checks the connectivity to RMS.

Metric	Alert item	Description
	regiondb_connectivity	Checks the connectivity to the houyiregiondb database.
	houyi_connectivity	Checks the connectivity to the Server Controller database.
	tair_connectivity	Checks the connectivity to TAIR.
check_zookeeper_work	status	Checks the operating state of the Zookeeper process on the Server Controller.
check_regionmaster_errorlog	errorlog_for_db	Checks whether the SQL statements are properly executed.
	check_regionmaster_errorlog	
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.	-

2.1.8.5 RMS

This topic describes the metrics of RMS.

Table 2-10: Metric description

Metric	Alert item	Description
check_rms_proc	Checks the process status, CPU utilization, and memory usage of RMS.	-
check_rabbitmq_proc	Checks the process status, CPU utilization, and memory usage of the rabbitmq cluster.	-

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.
check_rabbitmq_queues	Checks whether messages are accumulated.	If messages are accumulated, it will also check for the cause.
	Check whether there are consumers.	If there are no consumers, check whether Regionmaster and APIServer are operating normally. If they are operating normally, check whether there is a problem with the rabbitmq cluster.

2.1.8.6 PYNC

This topic describes the metrics that are monitored for PYNC.

Table 2-11: Metric description

Metric	Alert item	Description
check_vm_start_failed	Checks the causes of a VM startup fault.	You do not need to handle it immediately. It is typically caused by custom images.
check_pync	Checks the CPU utilization and memory usage of PYNC.	-
	PYNC has too many open file handles.	-
	PYNC process count.	PYNC must have four processes.

Metric	Alert item	Description
	It has been long since pyncVmMonitor.LOG was last updated at \${pync_monitor_log_last_updated}.	<p>Checks for reasons why a log has not updated for a long period of time, such as:</p> <ul style="list-style-type: none"> • Whether a PYNC process has encountered a problem. • Whether the NC is running a key process called Uninterruptible Sleep.

2.1.8.7 Zookeeper

This topic describes the metrics of Zookeeper.

Table 2-12: Metric description

Metric	Alert item	Description
check_zookeeper_proc	proc	The process does not exist.
		The memory usage or CPU utilization is too high.

2.1.8.8 AG

This topic describes the metrics of AGs.

Table 2-13: Metric description

Metric	Alert item	Description
disk_usage	apsara_90	/apsara disk usage.
	homeadmin_90	Usage of /home/admin.
check_system_ag	mem_85	Memory usage.
	cpu_98	CPU utilization.
	df_98	Disk usage of the root directory.
check_ag_disk_usage	check_ag_disk_usage	Disk usage.

Metric	Alert item	Description
check_nc_d own_new	check_recover_failed	Checks the causes of a VM migration fault. Possible causes include: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • If a hardware fault occurs , you must perform operations and maintenance to resolve the fault. • If no hardware fault is detected, restore the NC and change its state to locked.
check_ag_f htd_new	Checks whether the FHT downtime migration tool, mostly used by local disks, is operating normally.	If the tool does not exist, download the FHT downtime migration tool.

2.1.8.9 Server groups

This topic describes the metrics that are monitored for server groups.

Table 2-14: Metric description

Metric	Alert item	Description
check_pync	pync_mem	Monitors the memory usage of PYNC.
	pync_cpu	Monitors the CPU utilization of PYNC.
	pync_nofile	Monitors the number of PYNC handles.

Metric	Alert item	Description
	pync_nproc	Monitors the number of PYNC processes.
	pync_monitor_log_not_updated	Monitors the status of PYNC scheduled tasks.

2.1.9 Inspection

2.1.9.1 Overview

ECS inspection includes cluster basic health inspection and cluster resources inspection.

2.1.9.2 Cluster basic health inspection

2.1.9.2.1 Overview

Cluster basic health inspection includes monitoring inspection, inspection of basic software package versions, and basic public resources inspection.

2.1.9.2.2 Monitoring inspection

This topic describes basic monitoring inspections and connectivity monitoring inspections.

2.1.9.2.3 Inspection of basic software package versions

This topic describes the version inspections of Server Controller components, Apsara system, virtualization packages, and basic service packages.

2.1.9.2.4 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/put meta $path` command to check whether the ISO exists in the files of Apsara Distributed File System.
- **oss_info:** the path on the local OSS disk where the ISO file is stored. To search for this path, you must provide relevant information to OSS support engineers for inspection.

Basic image inspection

- **Run the following command to check the state of a basic image in the database:**

```
houyiregiondb
mysql>select image_no,status,visibility,platform,
region_no from image;
```

- **Check whether the basic image is usable. You can call the `create_instance` API to use relevant images to create a VM and manually check whether the VM can operate normally.**

2.1.9.3 Cluster resource inspection

2.1.9.3.1 Overview

Cluster resource inspection includes cluster inventory inspection and VM inspection.

2.1.9.3.2 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_resource,nc where nc.cluster_id=$id and nc.biz_status='free' and nc.id=nc_resource.id;
```

If the current cluster contains a relatively large VM, ensure that the cluster has enough free resources to handle the VM, as well as an available host with sufficient resources for backup. This host will be the migration destination of the large VM in case the current host goes down. Otherwise, the large VM cannot be migrated when its host goes down, and you will have to either use hot migration to transfer resources or release redundant VMs in the cluster.

NC state inspection

NC state inspection mainly checks whether the state of a host is normal in the database and Apsara Infrastructure Management Framework.

- A host can be in one of the following states in Apsara Infrastructure Management Framework:
 - **Good:** indicates that the host is in a normal working state.
 - **Error:** indicates that the host has an active monitoring alert.
 - **Probation:** indicates that the host is in the probationary period and may fail.
 - **OS _error:** indicates that the host has failed and is being cloned.
 - **Hw_error:** indicates that the hardware of a host has failed and is being repaired.
 - **OS _probation:** indicates the host is recovering from a fault or hardware failure and is in a probationary period. If the host recovers within the probationary period, the state will change to probation. If the host fails to recover within the probationary period (an error is reported), the state will change to OS _error.



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 *Table 2-15: Mappings of host states between the ECS database and Apsara Infrastructure Management Framework*.

Table 2-15: Mappings of host states between the ECS database and Apsara Infrastructure Management Framework

Host states in ECS database	Cluster	Host state	Scenario
mlock	Non-default cluster	Unstable	A host that goes online is immediately and proactively locked.
locked	Non-default cluster	Unstable	An NC needs to be unlocked.
free	Non-default cluster	Stable	A host operates normally.
nc_down	Non-default cluster	Unstable	A host operates normally or is in downtime.
offline	Default cluster	Unstable	A host goes offline from business attributes.

2.1.9.3.3 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.

2.2 Container Service

2.2.1 Components and features

2.2.1.1 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`

2.2.1.2 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.

2.2.1.3 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 left-side 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.

2.2.2 System restart

2.2.2.1 Restart a control node

A container control node is a Docker container where a service, such as CosConsoleAliyunCom, Troopers, or Etc, is deployed. To restart a control node, perform the following operations:

Procedure

1. 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 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 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.

2.2.3 Security maintenance

2.2.3.1 Network security maintenance

This topic describes how to configure ECS security groups.

2.3 Auto Scaling (ESS)

2.3.1 Log on to Apsara Stack Operations

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

Prerequisites

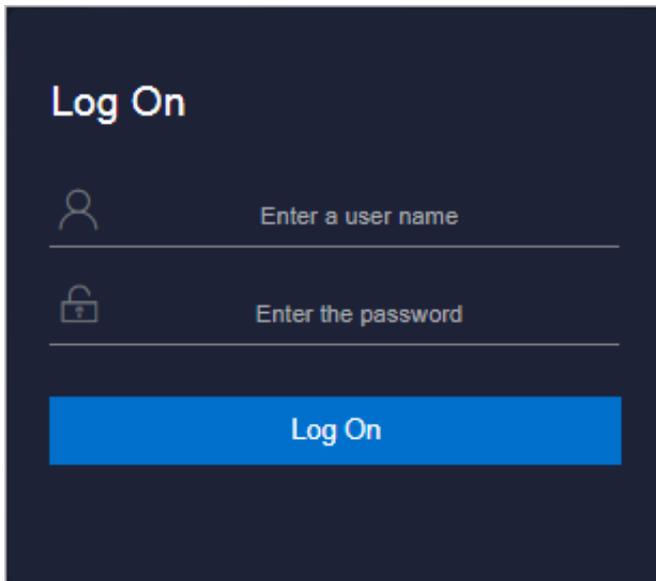
- ASO access address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

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.

Figure 2-3: Log on to ASO

**Note:**

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

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

2.3.2 Product resources and services

2.3.2.1 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.

2.3.2.2 Troubleshooting

This topic describes how to troubleshoot issues of product resources and services.

Prerequisites

When issues related to Business Foundation System occur, you can submit tickets on the [Alibaba Cloud Business Support Platform](#) and check related service status in the Apsara Infrastructure Management Framework console.

Procedure

1. **Submit a ticket.**

2. Check the status of services that depend on Business Foundation System in the Apsara Infrastructure Management Framework console.

If a service cannot be executed, it affects the running of ESS Business Foundation System. *Table 2-16: Failed services and their impacts* describes the details.

Table 2-16: Failed services and their impacts

Service	Impact
middleWare.dubbo	Deployment is affected. The service is unavailable.
middleWare.tair	Deployment is affected. The service is unavailable.
middleWare.metaq (message middleware)	Deployment is affected.
middleWare.zookeeper	Deployment is affected. The service is unavailable.
middleWare.jmenvDiamondVips	Deployment is affected, the Diamond configuration item cannot be obtained.
ram.ramService (RAM users)	The RAM-user service is unavailable.
webapp.pop (API gateway)	The OpenAPI service is unavailable.
ecs.yaochi (ECS Business Foundation System)	All ECS creation requests become invalid.
slb.yaochi (SLB Business Foundation System)	All SLB association requests become invalid.
rds.yaochi (RDS Business Foundation System)	All RDS association requests become invalid.
tianjimom (Monitoring System of Apsara Infrastructure Management Framework)	Some services are unavailable.

2.3.3 Inspection

2.3.3.1 Overview

ESS inspection monitors the basic health conditions of the clusters.

The basic health conditions inspected include the following aspects:

- *Monitoring inspection*

- *Basic software package version inspection*

2.3.3.2 Monitoring inspection

This topic describes basic monitoring and connectivity monitoring inspection.

2.3.3.3 Basic software package version inspection

Version inspection for trigger, coordinator, worker, and base services.

2.4 Object Storage Service (OSS)

2.4.1 Log on to the Apsara Stack Operations console

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

Prerequisites

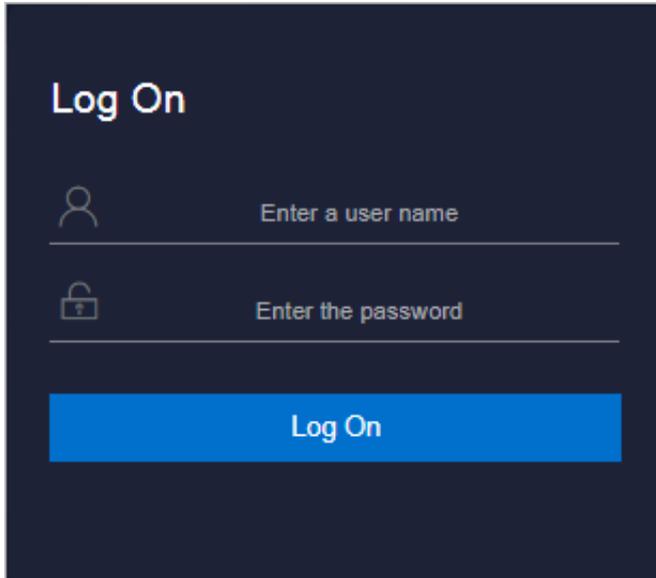
- ASO access address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

Procedure

1. Open the browser.

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

Figure 2-4: Log on to ASO



Note:

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

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

2.4.2 OSS operations and maintenance

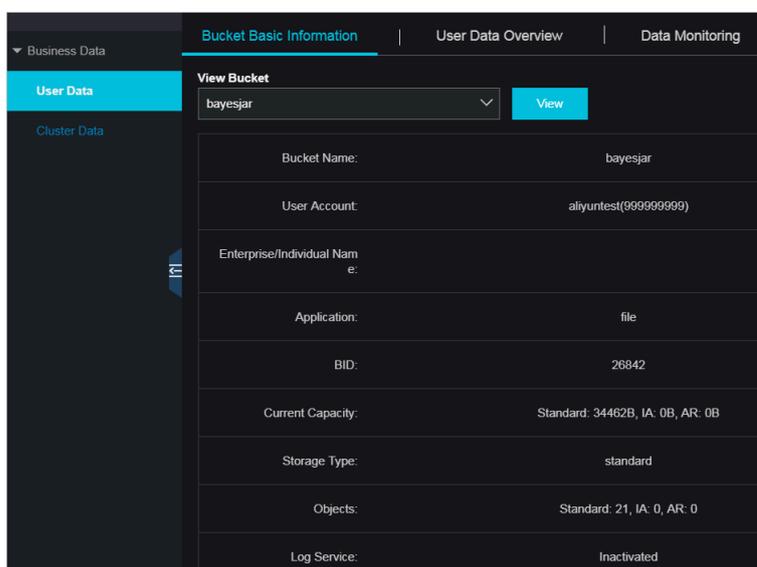
2.4.2.1 User data

2.4.2.1.1 Basic bucket information

You can query basic bucket information such as the cluster deployment location, configuration information, current capacity, and object count of a bucket. You can also view this information in a table.

Procedure

1. *Log on to the Apsara Stack Operations console.*
2. In the left-side navigation pane, choose **Products > OSS > User Data**.
3. On the **Bucket Basic Information** tab, select the bucket you want to view.
4. Click **View**, as shown in the following figure.



2.4.2.1.2 User data overview

You can query data statistics and trends, including resource usage and basic attributes of resources by UID, Alibaba Cloud Account, Bucket Name, or Bucket MD5.

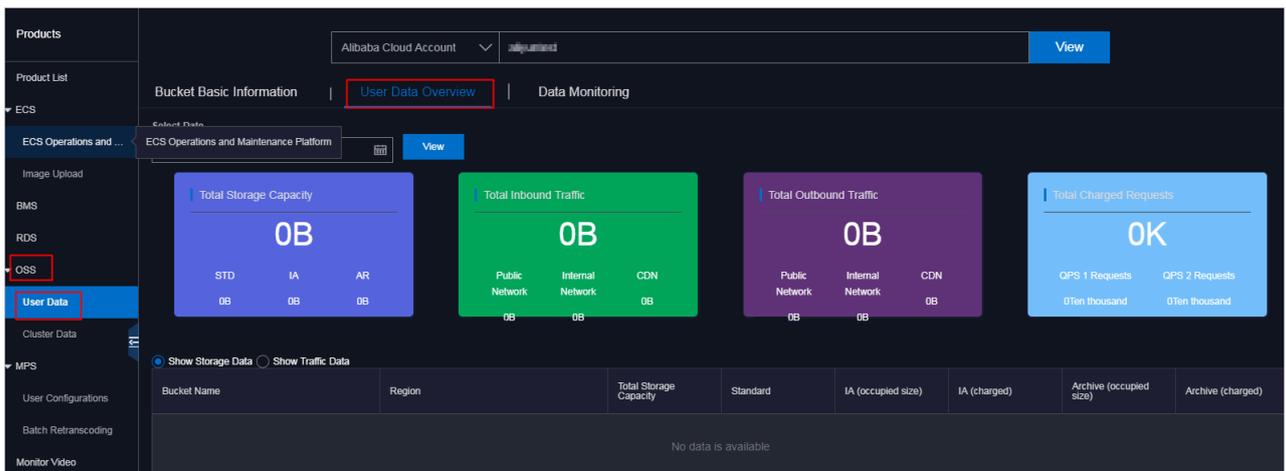
Context

The **User Data Overview** tab is displayed only when you search by UID or Alibaba Cloud account. On the **User Data Overview** tab, you can specify a date to view total usage of various resources in all buckets owned by the user account.

You can collect resource statistics by total storage capacity, total inbound or outbound traffic through the public network, internal network, or CDN, or total charged requests.

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.
4. Set Date. Click **OK**. Click **View**, as shown in the following figure.



2.4.2.1.3 Data monitoring

This topic describes how to monitor OSS data in the Apsara Stack Operations console.

Context

You can query resource running statuses and usage such as the storage capacity, traffic, SLA, HTTP status, latency, QPS, and image processing capacity by UID, Alibaba Cloud Account, Bucket Name, or Bucket MD5. You can also query the resource usage and trends based on a specified time range.

Procedure

1. [Log on to the Apsara Stack Operations console.](#)
2. In the left-side navigation pane, choose **Products > OSS > User Data**.
3. On the **Data Monitoring** tab, set **Bucket Name**, **Specify Time Range**, and **Monitoring Items**.



Note:

Metric descriptions:

- **SLA:** indicates the service level availability metric for OSS. Formula: $SLA = \frac{\text{Non-5xx request count per 10s or hour}}{\text{Total valid request count}} \times 100\%$.
- **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.

Figure 2-5: Data monitoring 1



- Move the pointer over the trend chart to display data at a specific point in time.

Figure 2-6: Data monitoring 2



2.4.2.2 Cluster data

2.4.2.2.1 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.

Cluster Data

Inventory Monitoring | Bucket Statistics | Object Statistics | Data Monitoring | Resource Usage Ranking

Report Type: Storage Inventory

Data Dimension: Apsara Distributed File System Data Statistical Time: 01/02/2019 View

Sampling Time : 01/02/2019, 18:14:18 Refresh

Region	Cluster	Total Capacity (TB)	Used Capacity (TB)	Unused Capacity (TB)	Utilization	Data Increment (TB)			Actions
						D	W	M	
cn-qingdao-e	oss-hybridcluster-20181021-25a6	53.15	5.11	48.04	9.61%	45.5	14.38%	0.12	1.29 4.27

Remarks :

- Remaining days of peak increment is calculated based on 90% of the cluster storage;
- The data is green when Apsara Distributed File System utilization is 70%–85%, yellow when the utilization is over 85%, and red when Apsara Distributed File System expires in 30 days or the physical space of Apsara Distributed File System is two times larger than the OSS logical space.

Aside from basic cluster information such as the cluster name and region, you can also view metrics based on the following dimensions:

- **Apsara Distributed File System Data:** includes the actual total capacity for storage (including the total capacity for multiple data backups), used capacity, remaining capacity (available), usage, and backup ratio.
- **Metric Data:** includes the bucket storage used by users who use ECS instances and other instances.
- **KV Data:** includes the logic KV data, KV data in the recycle bin, and data increment (by day, week, or month).

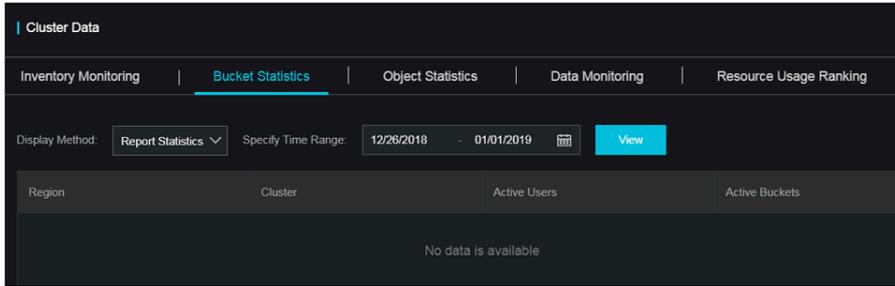
2.4.2.2.2 Bucket statistics

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.



- If you select Report, specify the time range.
- You can select Current Overall Statistics to query statistics of last hour.
- If you select Growth Trend, you can specify a time range of seven days, 30 days, three months, six months, or one year.

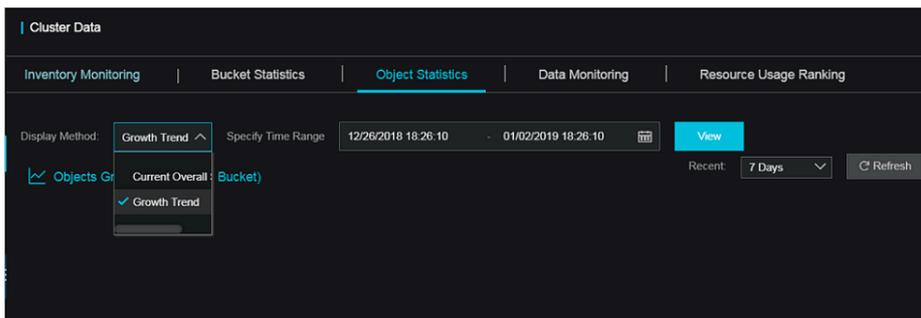
4. Click View.

2.4.2.2.3 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.



- You can select Current Overall Statistics to query statistics of last hour.
- If you select Growth Trend, you can specify a time range of seven days, 30 days, three months, six months, or one year.

4. Click View.

2.4.2.2.4 Data monitoring

This topic describes how to collect statistics for each metric by cluster.

Context

Cluster data metrics are similar to user data metrics except that the object of cluster data metrics is the data collected by cluster.

Procedure

1. [Log on to the Apsara Stack Operations console](#).
2. In the left-side navigation pane, choose **Products > OSS > Cluster Data**.
3. On the **Data Monitoring** tab, set **Monitoring Items** and **Specify Time Range**. Click **View**.



Note:

Metric descriptions:

- **SLA:** indicates the service level availability metric for OSS. Formula: $SLA = \frac{\text{Non-5xx request count per 10s or hour}}{\text{Total valid request count}} \times 100\%$.
- **Traffic:** collects statistics for the inbound and outbound traffic through the public network, internal network, and CDN and inbound and outbound synchronization traffic.
- **QPS:** collects statistics for the charged requests, including requests sent through the following API operations: CopyObject, GetObject, PutObject, UploadPart, PostObject, AppendObject, HeadObject, and GetObjectInfo.
- **Latency:** collects latency statistics for API operations such as PutObject, GetObject, and UploadPart as well as the maximum latency.
- **HTTP Status:** collects statistics for the percentages of the numbers of 5xx, 403, 404, 499, 4xx_others, 2xx, and 3xx status codes out of total requests.
- **Storage Capacity:** collects statistics for the storage capacity of standard, Infrequent Access (IA), and archive buckets and their increments.

4. Move the pointer over the trend chart to display data at a specific point in time.

Figure 2-7: Data monitoring 1



Metric descriptions:

- **SLA:** indicates the service level availability metric for OSS. Formula: $SLA = \frac{\text{Non-5xx request count per 10s or hour}}{\text{Total valid request count}} \times 100\%$.
- **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.

The following example describes typical operations on the data monitoring trend chart:

- If you query data monitoring information by user, you can click the bucket name in the trend chart to show or hide the curve.

Figure 2-8: Data monitoring 2



- Move the pointer over the trend chart to display data at a specific point in time.

Figure 2-9: Data monitoring 2



2.4.2.2.5 Resource usage rankings

This topic describes how to collect usage of resources by cluster. This way, administrators can monitor users that consume more resources.

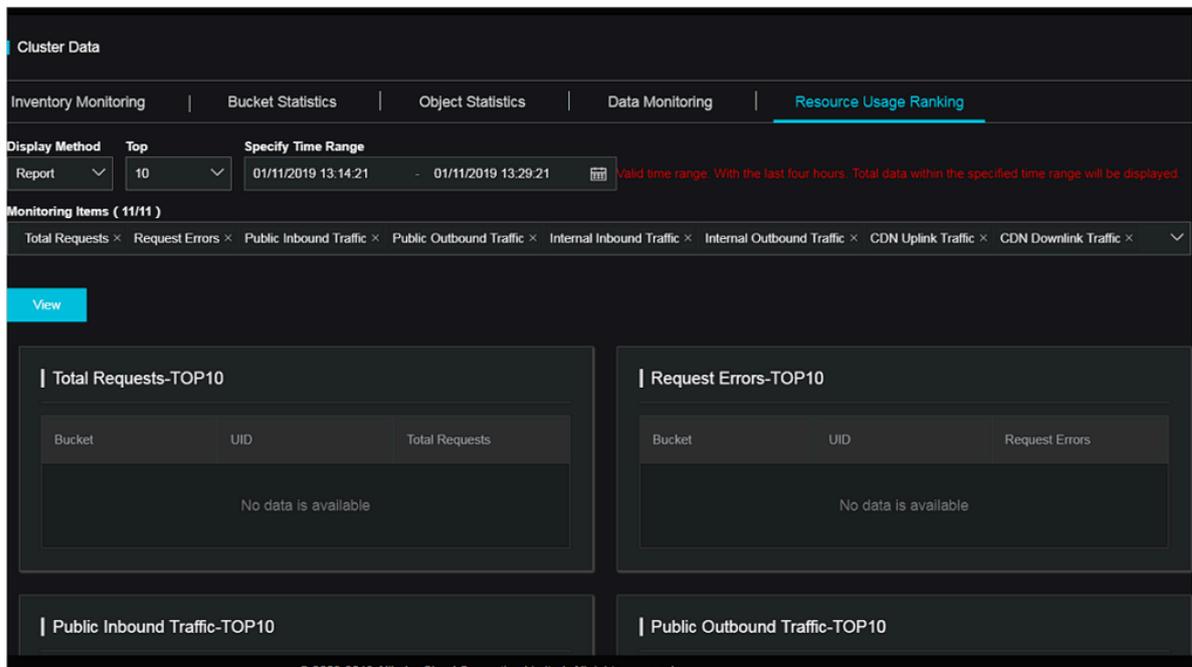
Context

Data resources can be ranked based on the following metrics:

- Total Requests
- Request Errors
- Public Inbound Traffic and Public Outbound Traffic
- Internal Inbound Traffic and Internal Outbound Traffic
- CDN Uplink Traffic and CDN Downlink Traffic
- Storage Capacity, Storage Increment, and Storage Decrement

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** drop-down list. Select a number from the **Top** drop-down list. Set **Specify Time Range** and **Monitoring Items** to view resource usage.



- In report mode, you can view the top 10, 30, or 50 buckets by resource usage.
- In trend mode, you can view the top 10 buckets by resource usage.

4. Click View.

2.4.3 Use of tools

2.4.3.1 Typical commands supported by tsar

You can use tsar to perform operations and maintenance on OSS. This topic describes typical commands supported by tsar.

tsar allows you to run the following commands:

- View help details of tsar

Command: `tsar -help`

- View the NGINX operation data of each minute from the past two days

Command: `tsar -n 2 -i 1 -nginx`

In this command, `-n 2` indicates the data generated in the past two days. `-i 1` indicates one result record generated each minute.

- View the tsar load status and operation data of each minute from the past two days

Command: `tsar --load -n 2 -i 1`

2.4.3.2 Configure tsar for statistic collection

You can configure tsar to collect data generated when NGINX runs.

Run the following command to configure tsar for statistic collection:

```
cat /etc/tsar/tsar.conf |grep nginx
```

The following figure shows that the status of `mod_nginx` is on.

```
admin:~# /home/admin
$cat /etc/tsar/tsar.conf |grep nginx
mod_nginx on ← Ensure that this item is in the on state.
output_stdio_mod mod_swap,mod_partition,mod_cpu,mod_mem,mod_lvs,mod_haproxy,mod_traffic,mod_squid,mod_load,mod_tcp,mod_udp,
mod_tcpx,mod_apache,mod_pcs,mod_io,mod_percpu,mod_nginx,mod_tcprt
```

2.5 Table Store

2.5.1 Table Store Operations and Maintenance System

2.5.1.1 Overview

Table Store Operations and Maintenance System helps locate problems during O&M and notifies users of the current running status of their services. Appropriate use

of Table Store Operations and Maintenance System can significantly improve O&M efficiency.

The endpoint of Table Store Operations and Maintenance System is in the format of "chiji.ots.{\$global:intranet-domain}."

Table Store Operations and Maintenance System consists of the following modules: User Data, Cluster Management, Inspection Center, Monitoring Center, System Management, and Platform Audit. These modules provide comprehensive O&M functions to meet different requirements.

2.5.1.2 User data

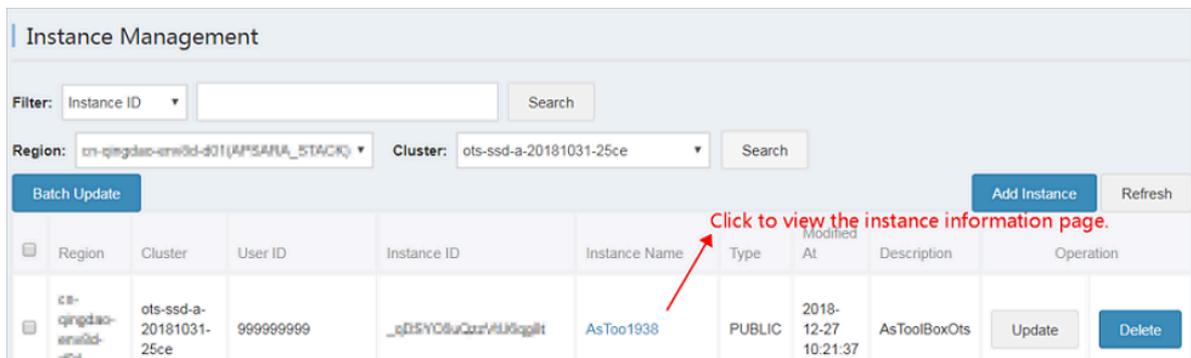
2.5.1.2.1 Instance management

You can obtain instance details through the cluster instance list, specified query conditions, and instance meta information.

Function description

- Specify a region and a cluster name to obtain instances.

You can specify a region and a cluster to view the instances, and the basic information of each instance in the specified cluster.



The screenshot shows the 'Instance Management' interface. At the top, there are filters for 'Instance ID' and 'Search'. Below that, there are dropdown menus for 'Region' (cn-qingdao-arcid-d01(APSARNA_STACK)) and 'Cluster' (ots-ssd-a-20181031-25ce), with a 'Search' button. There are also buttons for 'Batch Update', 'Add Instance', and 'Refresh'. The main part of the interface is a table with the following columns: Region, Cluster, User ID, Instance ID, Instance Name, Type, Modified At, Description, and Operation. The table contains one instance with the name 'AsToo1938'. A red arrow points to the 'Instance Name' column with the text 'Click to view the instance information page.'

Region	Cluster	User ID	Instance ID	Instance Name	Type	Modified At	Description	Operation
cn-qingdao-arcid-d01	ots-ssd-a-20181031-25ce	999999999	_cD5Y0duQzzr#U8qglt	AsToo1938	PUBLIC	2018-12-27 10:21:37	AsToolBoxOts	Update Delete

On the Instance Management page, you can:

- View the instances in the cluster.
- View instance descriptions.
- View the links to details of instances by clicking instance names.
- Update and delete an instance in the instance list.

- Search for instances based on specified conditions.

This page allows you to search for instances of all clusters in all regions based on the specified filtering conditions.

The screenshot shows the 'Instance Management' interface. At the top, there is a search bar with a dropdown menu set to 'Instance ID' and a text input field containing the instance ID '_qFmUwYyICdK7uqa2A'. Below this, there are dropdown menus for 'Region' and 'Cluster', both currently set to '--'. Action buttons include 'Batch Update', 'Add Instance', and 'Refresh'. A table displays the instance details:

Region	Cluster	User ID	Instance ID	Instance Name	Type	Modified At	Description	Operation
m-jagdoe-env64-d01	5ery-a-25ec	1385544530890241	_qFmUwYyICdK7uqa2A	asootsins	INTERNAL	2018-12-11 14:32:30		Update Delete

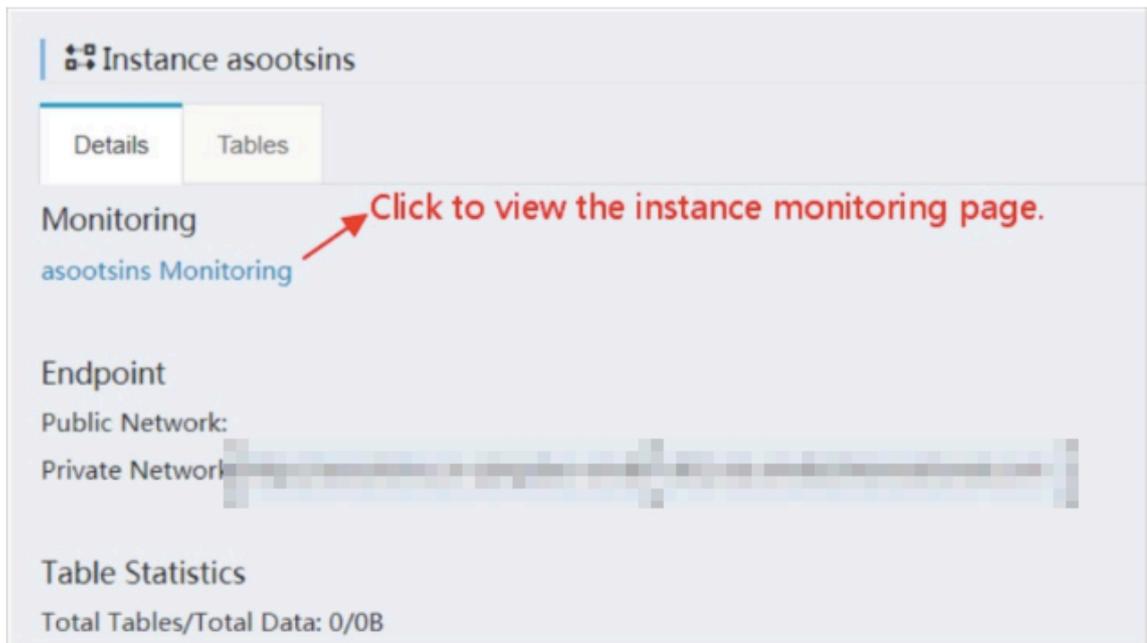
The available filtering conditions include:

- Instance ID
- Instance name
- User ID
- Apsara Stack account

- **View instance details.**

- **Instance overview**

Click the otssmoke96 instance to go to the Details tab. This tab provides detailed information about the instance, such as the instance monitoring link, intranet and Internet URLs, and statistics on tables in the instance.



- **Table information**

Click the Tables tab to view table information such as the max version, TTL, read CU, write CU, and timestamp.

The screenshot shows the 'Instance odps' interface with the 'Tables' tab selected. It displays a table with the following columns: Table Name, Max Version, TTL(s), Read CU, Write CU, Partitions, Data Size, Pangu Data Size, and Timestamp. Two tables are listed: 'ODPS_META_X_META_HISTORY' and 'ODPS_META_X_CHANGE_LOGS'. A red arrow points to the 'ODPS_META_X_META_HISTORY' table, with the text 'Click to view the table information page.' written in red next to it.

Table Name	Max Version	TTL(s)	Read CU	Write CU	Partitions	Data Size	Pangu Data Size	Timestamp
ODPS_META_X_META_HISTORY	1	-1	0	0	1	0B	59.4MB	2019-02-02 11:00:13
ODPS_META_X_CHANGE_LOGS	1	-1	0	0	1	0B	2720.8KB	2019-02-02 11:00:13

- **View table details.**

- **Details**

On the Tables tab, click the test_base_monitor table. On the Details tab, you can view the link to the monitoring data for this table, as well as the summary information such as the number of partitions and table data size.

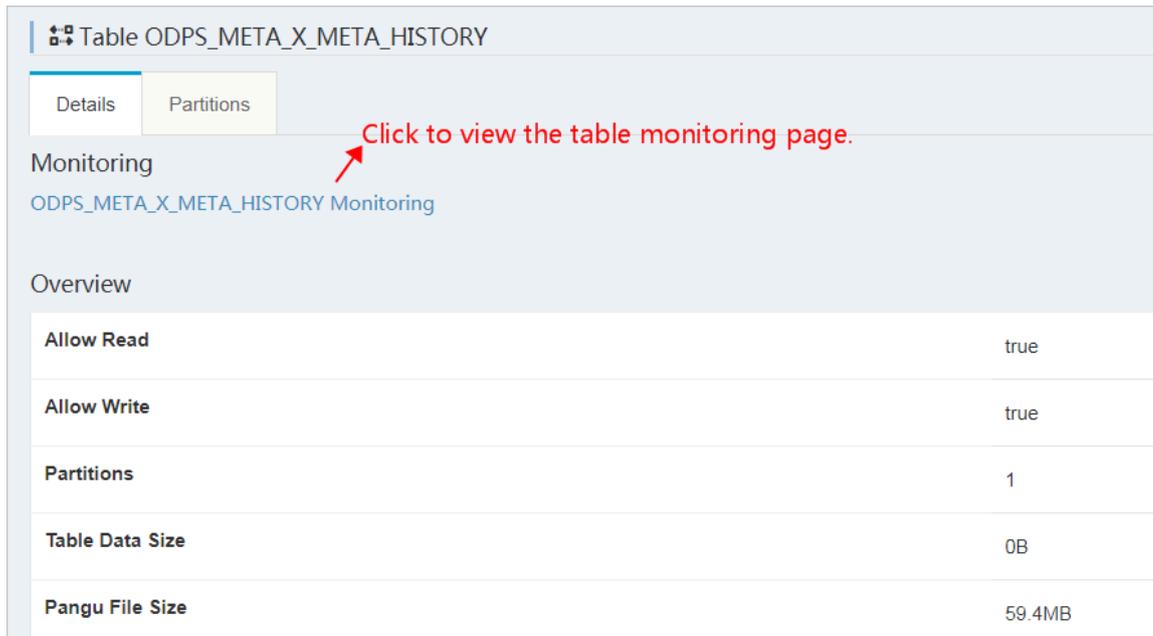


Table ODPS_META_X_META_HISTORY

Details | Partitions

Monitoring
[ODPS_META_X_META_HISTORY Monitoring](#)

Overview

Allow Read	true
Allow Write	true
Partitions	1
Table Data Size	0B
Pangu File Size	59.4MB

- **Partitions**

You can obtain the basic information of a partition, such as the partition ID and worker information. You can also specify filtering conditions to filter the partitions that meet your requirements.

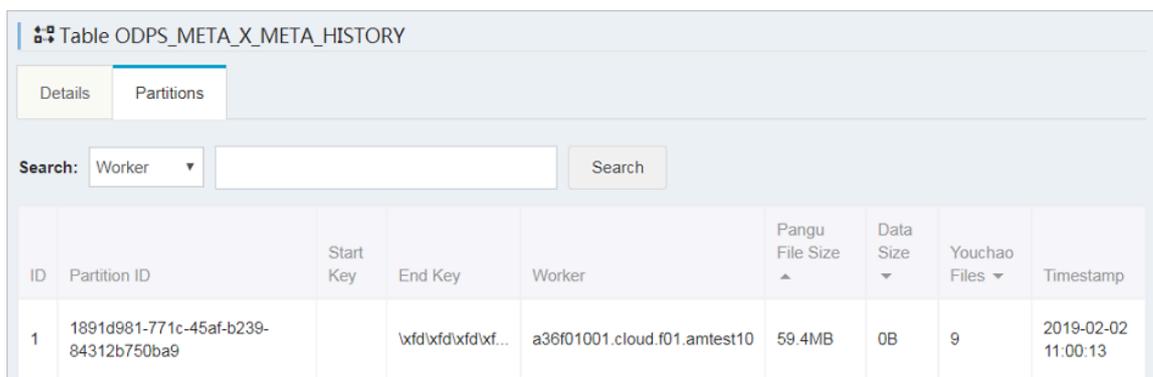


Table ODPS_META_X_META_HISTORY

Details | Partitions

Search: Worker Search

ID	Partition ID	Start Key	End Key	Worker	Pangu File Size	Data Size	Youchao Files	Timestamp
1	1891d981-771c-45af-b239-84312b750ba9		\xfd\xfd\xfd\xfd...	a36f01001.cloud.f01.amtest10	59.4MB	0B	9	2019-02-02 11:00:13

The available filtering conditions include:

- **Worker** (For more information, see the value in the Worker column.)
- **Partition ID**

2.5.1.3 Cluster management

2.5.1.3.1 Cluster information

You can obtain cluster information through cluster searches, cluster usage, and top requests.

Function description

- **Clusters**

Status	Cluster	Region	Storage Type	Operation
using	ots-hy-a-20181217-2e46	cn-qingdao-env8...	HYBRID	Delete
using	ots-ssd-a-20181031-25ce	cn-qingdao-env8...	SSD	Delete
using	tianji-a-25ee	cn-qingdao-env8...	HYBRID	Delete

Select All or specify a specific region from which to obtain clusters. The functions are as follows:

- **OCM cluster synchronization:** If you deploy an OCM service in each region of Table Store, the OCM service contains all cluster information of that region . This function synchronizes OCM clusters with their respective regions in Table Store Operations and Maintenance System to obtain all clusters in the regions.
- **Cluster deletion:** You can use this function to remove a cluster from Table Store Operations and Maintenance System after you confirm that the cluster is offline.

• Cluster details

Cluster Information				
Region: All		OCM Cluster Synchronization		Refresh
Status	Cluster	Region	Storage Type	Operation
using	ots-hy-a-20181217-2e46	cn-qingdao-env8...	HYBRID	Delete
using	ots-ssd-a-20181031-25ce	cn-qingdao-env8...	SSD	Delete
using	tianji-a-25ee	cn-qingdao-env8...	HYBRID	Delete

Click to view the cluster information page.

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-20181217-2e46

Overview | Top | Resource Usage

*Region: cn-qingdao-env8-d01 (APSARA_STACK) *Cluster: ots-hy-a-20181217-2e46 Switch Cluster

Region Description	APSARA_STACK
Region	cn-qingdao-env8-d01
Cluster	ots-hy-a-20181217-2e46
Armory App	mock_armory
Gateway	mock_ag
Cluster Type	public

- Top: provides top request information by partition and table.

Cluster ots-hy-a-20181217-2e46

Overview | Top | Resource Usage

Click to view the information page of top requests.

Click to view the table information page.

Click to view the partition information page.

Click to view the instance information page.

Top Partitions by Pangu File Size			Top Partitions by Youchao Files		
Table Name	Partition ID	Pangu File Size	Table Name	Partition ID	Youchao Files

Top Tables by Pangu File Size			Top Tables by Youchao Files		
Instance Name	Table Name	Pangu File Size	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.)

Cluster `ots-hy-a-20181217-2e46`

Overview | Top | **Resource Usage** | [Click to manually collect resource usage information](#)

Collected At: ~

Check Result :

Storage Resource Usage

Total Disk Size	Total File Size	Recycle Bin Size	Table Size	Free Space	Disk Usage Ratio (%)
					%

Gap Size | Hosts Total/Master/OTSServer/SqWorker | Hybrid Deployment | Cluster Type | Scale-out Requirement

	///			
--	-----	--	--	--

OTSServer Resource Usage

Hosts	Failed Hosts	Avg/Max CPU Usage (%)	Increased CPU Cores	Avg/Max NetIn (MB/s)	Increased Hosts Due to Excessive NetIn	Avg/Max NetOut (MB/s)	Increased Hosts Due to Excessive NetOut
		/		/		/	

[Collect Data](#)

2.5.1.4 Inspection center

2.5.1.4.1 Abnormal resource usage

You can click **Abnormal Resource Usage** in the left-side navigation pane to locate all cluster abnormalities and their causes.

Function description

Abnormal Resource Usage									
Cluster Name	Abnormal Resource Usage								
	Date	Total Disk Size	Total File Size	Gap	Recycle Bin Size	Table Size	Free Space	Disk Usage Ratio (%)	Scale-out Requirement
Sample-25	2019-02-02	64.46TB	6.21TB	3.25TB	1.64TB	1.32TB	48.80TB	24.31%	<p>35.27GB/Day, Reach Safe Level in -1Days, Growth Rate:-35.27GB/Days</p> <p>35.28GB/Day, Reach Safe Level in -1Days, Growth Rate:-35.28GB/Days</p>

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.

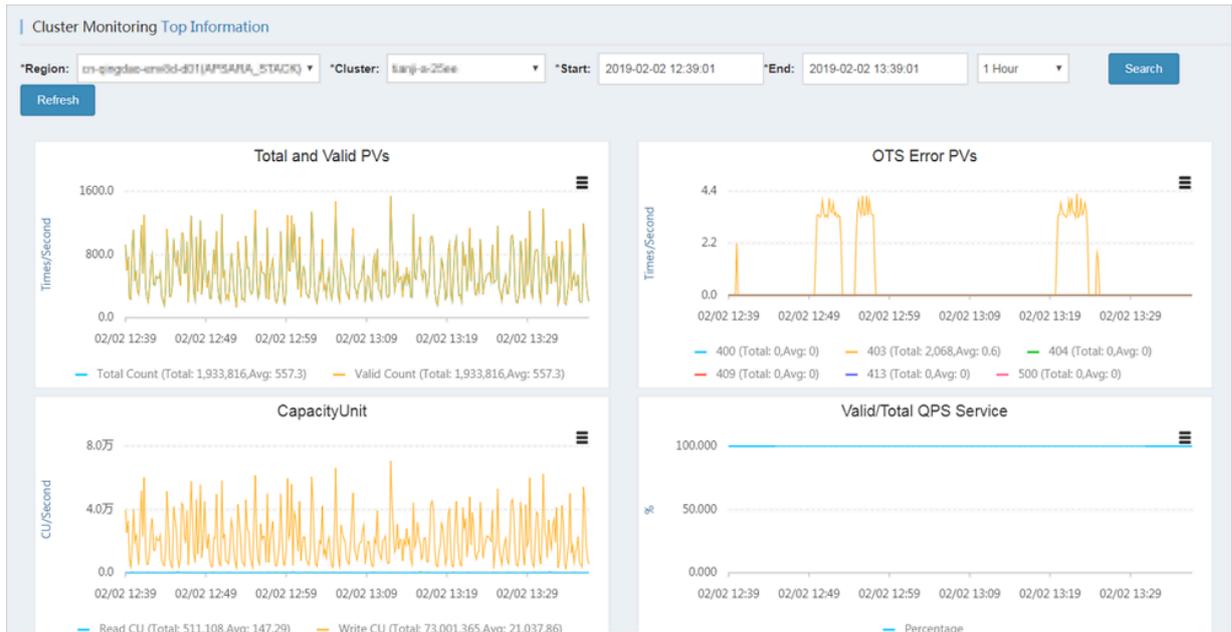
2.5.1.5 Monitoring center

2.5.1.5.1 Cluster monitoring

You can determine the service status of a cluster based on a series of metrics such as cluster-level monitoring information.

Function description

You can query the cluster service metrics within a specified time range, and determine whether a cluster service is healthy based on the metrics in the following dimensions.



2.5.1.5.2 Application monitoring

You can check the instance-level and table-level metrics to determine whether a service that belongs to a user is abnormal.

Function description

You can check the following metrics to determine whether a service for a specified user is in the healthy state.



Note:

The Instance field is required. Table and Operation fields are optional.



2.5.1.5.3 Top requests

You can view the top request distribution of clusters by monitoring level and dimension.

Function description

Four monitoring levels are supported for top requests: Instance, Instance-Operation, Instance-Table, and Instance-Table-Operation. You can view the top

request details of a cluster based on 13 different metrics, such as the total number of requests and the total number of rows.

Topic	Total Requests	Total Rows	Total Failed Rows	Public Uplink	Public Downlink	Internal Uplink	Internal Downlink	Read CPU	Write CPU	Total Latency Max Avg	SQLWorker Latency Max Avg	HTTP Status	SQL Status
{instanceName=metric...	1,643,542	73,033,406	0	0B	0B	19.3GB	1308.2MB	245,919	73,070,441	614,911 us 13,686 us	613,801 us 12,844 us	{200:1643542}	{0:73175642}
{instanceName=odps...	186,686	185,768	0	0B	0B	45.4MB	100.7MB	180,059	11,366	203,426 us 885 us	203,288 us 748 us	{200:186686}	{0:186686}

2.5.1.5.4 Request log search

You can search for a log entry based on a request ID to streamline problem investigation.

Function description

Query all log information about a request based on the request ID.

Host	Timestamp	File	Content
------	-----------	------	---------

2.5.1.6 System management

2.5.1.6.1 Manage tasks

You can maintain the back-end tasks in Table Store Operations and Maintenance System.

Function description

After Table Store Operations and Maintenance System is deployed in the Apsara Stack environment, the back-end tasks that collect usage statistics are automatically integrated. You can perform the following operations on the back-end tasks:

- **View task details such as the specific parameters and running time of each task.**

- Enable or disable a task.



Note:

Disabled tasks no longer run automatically.

- Run a task immediately.

The following figure shows the monitoring task details page. Based on the monitoring rules, the task collects usage statistics at 2:00 am every day.

Monitoring Task Details	
Task ID	1
Task Name	collect_water_level
Task Script	
Task Script Parameter	
Remote HTTP Task URL	http://10.68.163.205/ots/apsarastack/v1/inner/httptask/run
Cluster	
Host Role	
Monitoring Rule	0 0 2 * * ?
Task Status	1
Alert Receiver Employee ID	
DingTalk Group Chat Robot Webhook	
Task Type	4
Alert Method	0
Task Result Format	0

2.5.1.6.2 View tasks

You can view the execution status of back-end tasks and locate the causes of task exceptions.

The following figure shows the execution status of back-end tasks in Table Store Operations and Maintenance System. You can view the tasks, which have either succeeded or failed.

Status	Name	Type	Started At	Ended At	Operation
Abnormal	collect_water_level	Remote HTTP	2019-02-02 06:00:00	2019-02-02 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-02-01 06:00:00	2019-02-01 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-31 06:00:00	2019-01-31 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-30 06:00:00	2019-01-30 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-29 06:00:00	2019-01-29 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-28 06:00:00	2019-01-28 06:00:10	View All View Exceptions
Abnormal	collect_water_level	Remote HTTP	2019-01-27 06:00:00	2019-01-27 06:00:10	View All View Exceptions

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 result

total 1 count, 0 execute success, /1 execute fail, 1 execute warning

Executelp	StartTime	EndTime	TaskResult	Warning	IsSuccess
HTTP	Feb 2, 2019 2:00:00 AM	Feb 2, 2019 2:00:10 AM	"env: APSARA_STACK, inner task collect water level fail: Trigger collect water level fail, cluster list: [ots-hy-a-20181217-2e46, ots-ssd-a-20181031-25ce]"	env: APSARA_STACK, inner task collect water level fail: Trigger collect water level fail, cluster list: [ots-hy-a-20181217-2e46, ots-ssd-a-20181031-25ce]	fail

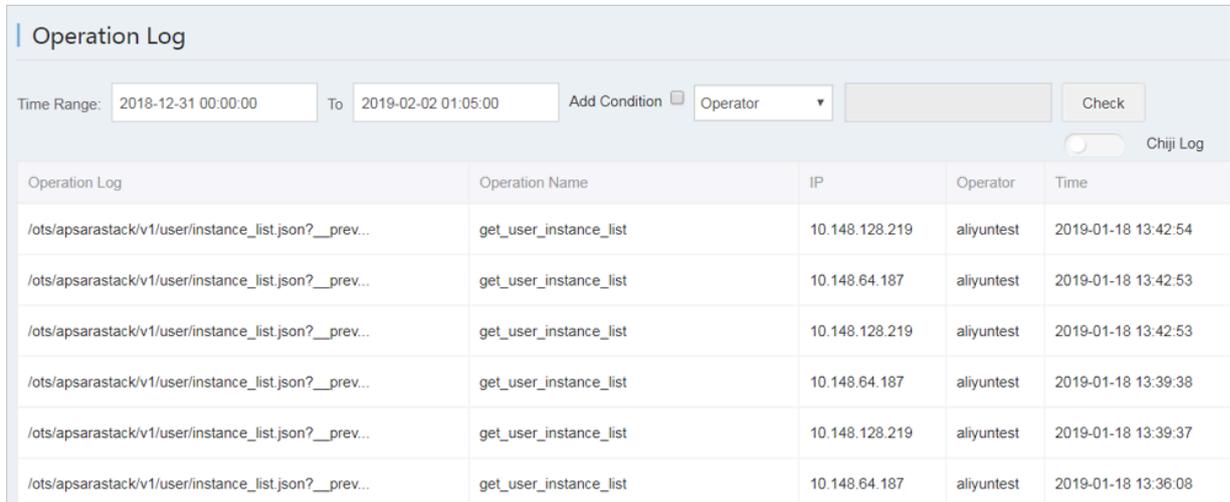
2.5.1.7 Platform audit

2.5.1.7.1 Operation logs

You can view the management and control operation logs of Table Store Operations and Maintenance System.

Function description

The Operation Log page provides the operation logs of Table Store Operations and Maintenance System. You can query audit records generated within a specified time range and filter the records as required. This helps management personnel obtain information about the platform status.



The screenshot shows the 'Operation Log' interface. At the top, there is a search bar with 'Time Range' set from '2018-12-31 00:00:00' to '2019-02-02 01:05:00'. There is an 'Add Condition' button and a dropdown menu for 'Operator'. A 'Check' button is also present. Below the search bar is a table with the following columns: Operation Log, Operation Name, IP, Operator, and Time. The table contains six rows of log entries.

Operation Log	Operation Name	IP	Operator	Time
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.128.219	aliyuntest	2019-01-18 13:42:54
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.64.187	aliyuntest	2019-01-18 13:42:53
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.128.219	aliyuntest	2019-01-18 13:42:53
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.64.187	aliyuntest	2019-01-18 13:39:38
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.128.219	aliyuntest	2019-01-18 13:39:37
/ots/apsarastack/v1/user/instance_list.json?__prev...	get_user_instance_list	10.148.64.187	aliyuntest	2019-01-18 13:36:08

2.5.2 Cluster environments

Two environments are provided for Table Store: the internal environment for cloud services such as MaxCompute, Log Service, and StreamSQL, and the external environment deployed for users.

Some cloud services use both environments simultaneously. For example, metadata of StreamSQL is stored in the internal environment, but its dimension table data (user data) is stored in the external environment.

Table Store services include TableStoreOCM, TableStoreInner/TableStore, TableStorePortal, chiji, and TableStoreSqlInner/TableStoreSql.

- **TableStoreOCM:** the tool used to manage information about clusters, users, and instances
- **TableStoreInner/TableStore:** the Table Store data service node
- **TableStorePortal:** the back-end of the Table Store O&M platform
- **chiji:** the Table Store O&M platform frequently used for fault location
- **TableStoreSqlInner/TableStoreSql:** the Table Store back-end tool

2.5.3 System roles

- **TableStoreOCM**
 - **OCMInit**: the OCM initialization tool used to create tables and bind POP APIs
 - **OCM**: the service node of OCM
 - **ServiceTest**: the service test image of OCM
- **TableStoreInner/TableStore**
 - **InitCluster**: the process of adding cluster information to OCM, including the domain name and type of the cluster, as well as the pre-configured Table Store account information
 - **LogSearchAgent**: the Table Store log collection service node
 - **MeteringServer**: the Table Store metering node (only available in Table Store)
 - **MonitorAgent**: the data collection node of the Table Store Monitor system
 - **MonitorAgg**: the data aggregation node of the Table Store Monitor system
 - **OTSAAlertChecker**: the Table Store alarm service module
 - **OTSFrontServer**: the front-end server of Table Store, which can be NGINX, OTS Server, or Replication Server
 - **OTSServer**: the OTS front-end server
 - **OTSTEngine**: the NGINX service for OTS front-end servers
 - **PortalAgServer**: the back-end 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 alarms but does not provide actual services now
 - **TableStoreAdmin**: all O&M tools of Table Store, including the splitting and merging tools
- **TableStorePortal**
 - **PortalApiServer**: the back-end service for Table Store Operations and Maintenance System
- **TableStoreSqlInner/TableStoreSql**
 - **Tools**: the back-end tools for Table Store, such as sqlonline_console
 - **UpgradeSql**: the back-end hot upgrade tool for Table Store

2.5.4 Pre-partition a table

2.5.4.1 Pre-partitioning

When you create a table, Table Store automatically creates a partition for the table. This partition can be configured to automatically split based on the data size or data access load as your business develops. A table with only one partition may be unable to provide sufficient service capabilities during a stress test or data import. In this scenario, you must pre-partition the table.

Pre-partitioning rules

You can estimate the number of partitions required based on the standard size of 10 GB per partition. However, considering other factors such as the number of hosts and concurrent write operations by developers, we recommend that the 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 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 partition key format

- **The partition key requires an integer value.**

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_digit table name
```

- **The partition key starts with the lowercase MD5 code ([0-9,a-f]).**

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_hex_lower table name
```

- **The primary key starts with the uppercase MD5 code ([0-9,A-F]).**

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_hex_upper table name
```

- **The partition key is encoded with Base64 ([+/0-9,A-Z,a-z]).**

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_base64 table name
```

- **--only_plan: generates split points but does not split the table. --force: directly splits the table without manual confirmation.**

```
python2.7 split_merge.py split_table -n (number of partitions) --
key_digit --only_plan table name
```

Split a partition based on the existing data

```
python2.7 split_merge.py split_partition -n PART_COUNT (number of
partitions) partition_id
```



Note:

You can also use the preceding methods to partition a data table that already stores data.

2.5.4.2 View partitions

You can view the partitions of a data table in Table Store Operations and Maintenance System.

On the homepage of Table Store Operations and Maintenance System, choose User Data > Instance Management from the left-side navigation pane. On the Instance Management page that appears, set Region and Cluster. Click Search. Locate an instance and click the instance name. The Details tab appears. Click the Tables tab. On the Tables tab, click a table name. The Details tab appears. Click the Partitions 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 is the size of the raw data of an actual user. The data may not be the real-time data, because the data is updated after files are merged in the back-end of the system. The Apsara Distributed File System file size is the compressed data size. (The actual storage space is triple the file size because the data is stored in three copies.)

2.6 Apsara File Storage for HDFS

2.6.1 Overview

2.6.1.1 Overview of Apsara Distributed File System

Apsara Distributed File System is the distributed file system component of the Apsara system. It integrates disks in low reliability PC servers into a whole and provides secure, stable, and easy-to-use file storage capability to external systems.

The following table lists the common terms used in Apsara Distributed File System.

Term	Description
multi-master	<p>To prevent a single point of failure caused by a single master, Apsara Distributed File System deploys multiple masters. One master serves as the primary master to provide services externally, and other masters provide backup as secondary masters. Once the primary master fails to provide services externally, Apsara Distributed File System automatically switches services to a secondary master to guarantee service availability.</p> <p>We recommend you to configure an odd number of machines as masters because Apsara Distributed File System adopts a simplified Paxos algorithm. The recommended number of masters is five.</p>
chunkserver	<p>The Apsara Distributed File System chunkserver used to store user data.</p>

Term	Description
supervisor	The module controlling the O&M of Apsara Distributed File System. After Apsara Distributed File System V0.16 is deployed in Apsara Infrastructure Management Framework, changes to Apsara Distributed File System clusters , such as hot upgrade, CS enabling, CS disabling, must be approved by the supervisor. The supervisor serves as an interface between Apsara Distributed File System and Apsara Infrastructure Management Framework. It provides approval and alert services, and is the data source of Pangu Portal.
monitor	A module that monitors hardware and environment statuses of machines and reports alarms to Apsara Infrastructure Management Framework when an exception occurs.
Oplog	As a service, Apsara Distributed File System keeps records of all operations in the form of operation logs (Oplogs). Once a machine restarts after fault occurrence, the data in its memory can be restored by replaying all Oplogs.
checkpoint	To reduce the number of Oplogs, Apsara Distributed File System periodically dumps data from the memory. The dumped files are called checkpoints. After a machine restarts, the latest checkpoint is loaded into the memory . Then, instead of replaying all Oplogs , you only need to replay the Oplogs generated after the checkpoint to restore the data.

2.6.1.2 Overview of Apsara Infrastructure Management Framework

Apsara Infrastructure Management Framework is an automatic data center management system that manages the hardware lifecycles and various static resources (such as programs, configurations, operating system images, and data) in a data center.

Apsara Infrastructure Management Framework provides a set of universal version management, deployment and hot upgrade solutions for the applications and services of Alibaba Cloud Apsara Stack. It implements automatic O&M on Apsara Infrastructure Management Framework-based services in a large-scale distributed environment, greatly improving the O&M efficiency and system availability.

2.6.1.2.1 Terms of Apsara Infrastructure Management Framework

This topic describes several terms used in Apsara Infrastructure Management Framework.

cluster

A logical set of physical machines that provide services.

service

A service is the software that provides specific functions in Apsara Infrastructure Management Framework. Each cloud product is a service. The service name is globally unique. We recommended that you set a service name by combining lowercase letters and a prefix of a BU name, for example, aliyun.oss.

A service corresponds to one service package, which is a standard tar.gz file. The directory structure of a service package must comply with the service package specifications of Apsara Infrastructure Management Framework.

A service can be deployed on a group of hardware servers (a cluster), to provide the related service capabilities. For example, Pangu, Fuxi, and Nuwa are all services.

server role

A service can be divided by function into one or more server roles. A server role is an indivisible deployment unit and indicates a certain functional component of a service running on a hardware server. Deploying a server role onto a server

indicates that the server provides the corresponding function. Multiple server roles, for example, PanguMaster and TianjiClient, can be deployed on the same server.

We recommend that you use UpperCamelCase to name server roles, for example, PanguMaster. To support the multi-tenant mode, the full name of a server role contains the service name prefix as the namespace, for example, pangu.PanguMaster.

server role instance

The instance of a server role that is deployed in a cluster. A server role instance is expressed by <ServerRoleName>#[instanceNO], where ServerRoleName is the name of the server role, and instanceNO is the instance number, which can be a combination of letters and digits. A number of instances of different server roles can be deployed on one server in a cluster. For example, several versions of PanguLib can be deployed in one cluster. Different instances of one server role are identified with a number sign (#) and a suffix, for example, PanguLib#56 and PanguLib#57.

application

A process-level service component contained by a server role. Each application works independently. Application is the minimum unit for deployment and upgrade in Apsara Infrastructure Management Framework, and can be deployed on every server.

2.6.2 Configuration update

2.6.2.1 Overview

After configurations are modified in the Apsara Infrastructure Management Framework console, Apsara Infrastructure Management Framework pushes the modification to each machine corresponding to the server role. A process automatically detects the changes in the configuration file and updates the configuration.

The configuration update covers the startup parameters, flags, `apsara_log_conf.json`, and other custom configurations.

You can also update the configurations by setting the flag value (which takes effect immediately) in the memory through `puadmin`. Then modify the configuration in

Apsara Infrastructure Management Framework. You need to record the modification on to make sure the modification remains effective after process restart.



Note:

After you modify the startup parameters and some flag values that must take effect after the restart, the configuration will be pushed, but it will take effect only after process restart. Before updating the configuration, you need to know whether the modification requires process restart to take effect.

2.6.2.2 Procedure

This topic describes how to use the Apsara Infrastructure Management Framework console or puadmin to modify the service configurations of Apsara Distributed File System.

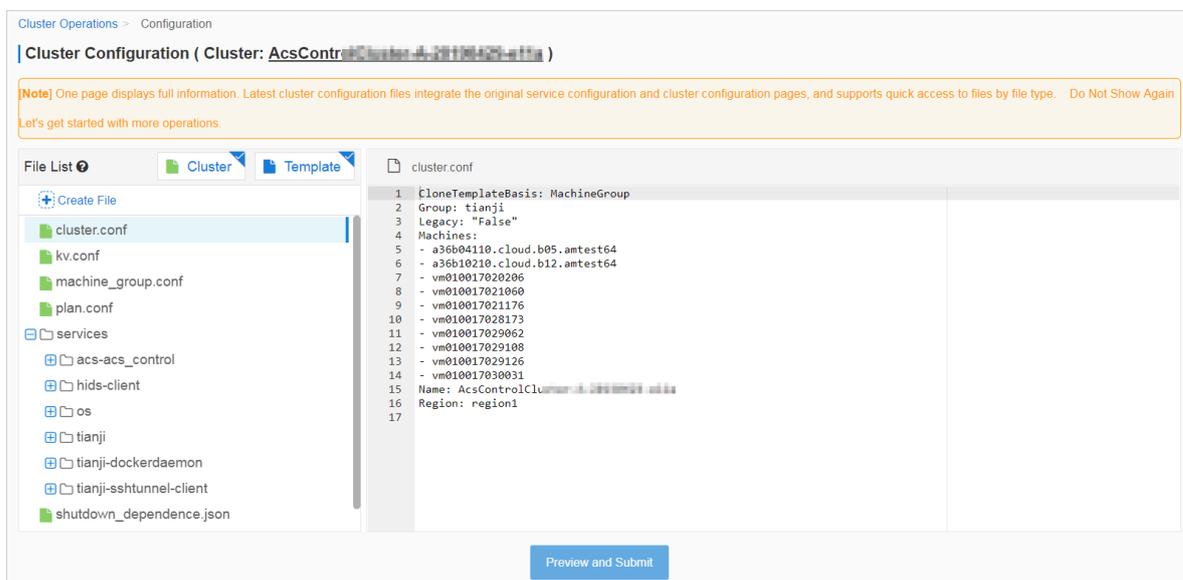
Modify configurations by using the Apsara Infrastructure Management Framework console

You can modify the service configurations of Apsara Distributed File System as follows in the Apsara Infrastructure Management Framework console.

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > Apsara Infrastructure Management Framework**.
3. Choose **Operations > Cluster Operations**. On the Cluster Operations page that appears, click **Edit** in the **Actions** column. The Cluster Configuration page is displayed as follows.

Cluster	Scale-Out/Scale-In	Abnormal Machine Count	Final Status of Normal Machines	Rolling	Actions
AcsControlCluster-A-20190429-acs	N/A	Good	The system reaches the final status.	Running History	Cluster Configuration Edit Management Monitoring
ads-A-20190429-e134	N/A	Good	The system reaches the final status.	Running History	Cluster Configuration Edit Management Monitoring
AliguardCluster-A-20190429-e-yundun-advance	N/A	Good	The system reaches the final status.	Running History	Cluster Configuration Edit Management Monitoring
amtest64-network	N/A	Good	The system reaches the final status.	Running History	Cluster Configuration Edit Management Monitoring
ansCluster-A-20190429-e190	N/A	Good	The system reaches the final status.	Running History	Cluster Configuration Edit Management Monitoring

4. After modifying or adding the corresponding file, click Preview and Submit, as shown in the following figure.



For more information about the directory structure of the configuration file, see [Configure the file structure](#).

Modify configurations using puadmin

puadmin is included in the PanguTools server role. You can use **puadmin** to set and view the flags.



Notice:

Flags configured through puadmin cannot remain persistent and will become invalid after the corresponding process is restarted.

- **Set flags**

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicateWindowSize 1572864 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicateWriteDataBlockSize 524288 -c
```

When changing the pangu_master flag, you must specify the -v parameter to indicate the volume to modify. The following example takes the v2 volume as an example. If the -v parameter is not specified, the default volume is configured.

```
/apsara/deploy/puadmin flag -set pangu_master_ReplicationTaskDestinationLengthLimit 6643777536 -v v2
```

- **Obtain flags**

The meaning of the `-v` parameter in the following example is the same as that in the flag setting command.

```
/apsara/deploy/puadmin flag -get pangu_chunkserver_ReplicateWindowSize  
-c
```

```
/apsara/deploy/puadmin flag -get pangu_chunkserver_ReplicateWriteDataBlockSize -c
```

```
/apsara/deploy/puadmin flag -get pangu_master_ReplicationTaskDestinationLengthLimit -v v2
```

**Notice:**

The preceding flags are test examples, where `-c` corresponds to chunkserver.

2.6.2.3 Configure the file structure

This topic describes how to configure the file structure in the Apsara Infrastructure Management Framework console.

When modifying configurations in the Apsara Infrastructure Management Framework console, you need to understand the directory structure of the configuration file. Taking Apsara Distributed File System as an example, the directory structure is as follows:

```

|-- pangu↵
|   |-- user↵
|       |-- pangu_chunkserver↵
|           |-- conf↵
|               |-- apsara_log_conf.json
|               |-- pangu_chunkserver.json
|               |-- pangu_chunkserver_flag.json
|       |-- pangu_master↵
|           |-- conf↵
|               |-- apsara_log_conf.json↵
|               |-- pangu_master.json↵
|               |-- pangu_master_flag.json↵
|               |-- pangu_master_volume.json
|               |-- pangu_master_volume_flag.json
|       |-- pangu_monitor↵
|           |-- conf↵
|               |-- pangu_monitor_items.json
|               |-- pangu_monitor.json
|               |-- pangu_monitor_flag.json
|       |-- pangu_supervisor↵
|           |-- conf↵
|               |-- apsara_log_conf.json↵
|               |-- pangu_supervisor.json↵
|               |-- pangu_supervisor_flag.json↵

```

The configuration files of `pangu_chunkserver`, `pangu_supervisor`, and `pangu_master` have the same structure.

- **Modify a flag in the `user/pangu_*/pangu_*_flag.json` file.**
- **Modify a startup parameter in the `user/pangu_*/pangu_*.json` file.**
- **Modify a log configuration file in the `user/pangu_*/conf/apsara_log_conf.json` file.**



Note:

Make sure that the directory structure is correct. Otherwise, unexpected results may occur.

The directory structure of `pangu_monitor` is slightly different from that of others.

- **The configuration file under `conf` describes the metric items, rather than log configurations.**

- The configuration file `pangu_monitor_items.json` under `monitor` is differentiated based on applications.

For the MaxCompute applications, `cs_load_usage` is set to generate a warning when the value reaches 400 and an error when the value reaches 500. To modify a metric value, select the corresponding application according to the value of `CLUSTER_TYPE` in the tag, copy the file to the Apsara Infrastructure Management Framework console, and modify the corresponding value.

You may need to add these configuration files manually because each template is different. `xx.json` describes the startup parameters and `xx_flag.json` describes the flags.



Note:

Among all configurations, a key is a string. Values must be enclosed in quotation marks. Values cannot contain spaces.

2.6.2.4 Make the configuration changes take effect

After a task is submitted, Apsara Infrastructure Management Framework pushes the modified configuration file and calls the `update_config` script provided by the application. The `update_config` script computes and saves the configuration differences under the `config_update` folder. The `pangu_master`, `pangu_chunkserver`, `pangu_supervisor`, and `pangu_monitor` periodically load and apply the configuration changes. Except for the startup parameters and flags that cannot take effect immediately, other configurations take effect immediately after pushing. For startup parameters and flags that cannot take effect immediately, execute the service upgrade process to make them take effect.

2.6.2.5 Overwrite configurations

The principle for configuration management is to save all configurations in the code, rather than making configuration changes directly on clusters. We recommend that you manage configurations in the Apsara Infrastructure Management Framework console only when configurations in the code cannot meet requirements. The configurations in the Apsara Infrastructure Management Framework console overwrite the configurations in the code.

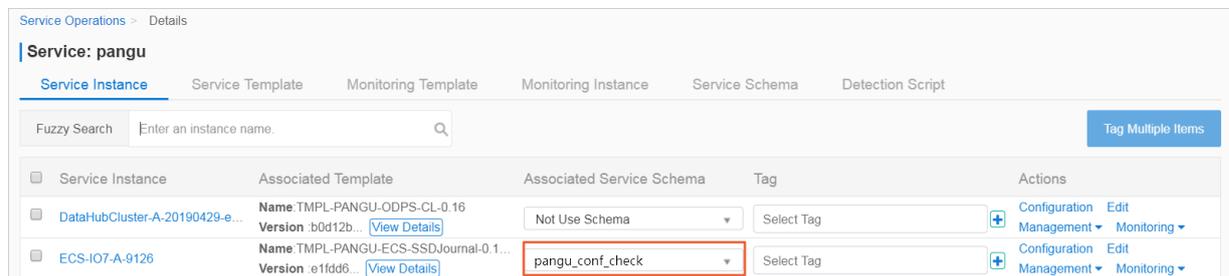
2.6.2.6 Configure validity checks

Configuration files are in JSON format. In Apsara Infrastructure Management Framework, a schema is defined to check the syntax of JSON files.

When a user tries to submit a JSON file that has syntactic errors, an error is returned and the submission fails. The user needs to modify the JSON file before submitting it again. This mechanism ensures all issued JSON files are valid. Perform the following steps to check the validity of a JSON file:

Choose **Operations > Service Operations**. On the Service Operations page that appears, click **Management** in the Actions column corresponding to **pangu**. On the page that appears, click the **Service Instance** tab. On the Service Instance tab, locate a cluster and then select a schema from the **Associated Service Schema** drop-down list.

In the following example, the cluster is correlated with the **pangu_conf_check** schema. To ensure the correctness of the configuration syntax, use the schema to correlate to the clusters of Apsara Distributed File System. If the schema is not correlated, the syntax check will not be performed.



2.6.3 Cluster operations and maintenance in Apsara Distributed File System

2.6.3.1 Set a global flag in Apsara Distributed File System

Apsara Distributed File System uses global flags to describe variables that can be modified externally and allows you to set and view global flags using **puadmin**. **puadmin** is an application of PanguTools. It is usually deployed on the AG and stored in the `/apsara/deploy/puadmin` directory.

Usage

The syntax format is as follows:

- `puadmin flag -set flag_name flag_value [option] [option]`

- `puadmin flag -get flag_name flag_value [option] [option]`

Examples:

- **Set a flag for the chunkserver:** `/apsara/deploy/puadmin flag -set pangu_chunkserver_xxx 300 -c`
- **Obtain a flag for the chunkserver:** `/apsara/deploy/puadmin flag -get pangu_chunkserver_xxx -c`
- **Set a flag for the master:** `/apsara/deploy/puadmin flag -set pangu_master_xxx 100000 -m` **or** `/apsara/deploy/puadmin flag -set pangu_master_xxx 100000`
- **Obtain a flag for a master:** `/apsara/deploy/puadmin flag -get pangu_master_xxx -m` **or** `/apsara/deploy/puadmin flag -get pangu_master_xxxx`

2.6.3.2 Manage files in Apsara Distributed File System

Apsara Distributed File System provides the `pu` tool for various file operations. For example, you can perform the following operations:

- **Create a folder named newdir:** `/apsara/deploy/pu mkdir pangu://localcluster/newdir/`
- **Create a folder named newdir:** `/apsara/deploy/pu rmdir pangu://localcluster/newdir/`
- **Upload the file named newfile to the newdir folder in Apsara Distributed File System:** `/apsara/deploy/pu cp newfile pangu://localcluster/newdir/`
- **Read the file named newfile from the newdir folder in Apsara Distributed File System to the local disk and name the file dstfile:** `/apsara/deploy/pu get pangu://localcluster/newdir/newfile dstfile`
- **Restore the newdir folder that has been deleted:** `/apsara/deploy pu restore pangu://localcluster/newdir/`



Note:

- Deleted files are stored in the recycle bin, that is, the *deleted* folder. The files and folders in this folder cannot be deleted by running the `pu rm` or `rmdir` command. To clear the recycle bin, run `/apsara/deploy/puadmin fs -`

`crb -f`. If the GC function is disabled in Apsara Distributed File System, the preceding commands do not take effect.

- You can run `./puadmin fs -quota pangu://localcluster/deleted/` to check the size of the `deleted` directory and determine the deletion progress.

2.6.3.3 Common puadmin commands

The puadmin tool is a common command line management tool of Apsara Distributed File System. It allows you to check the status of Apsara Distributed File System, modify its flags, and change its running status.

The common commands are as follows:

- Query the storage space statistics of Apsara Distributed File System: `puadmin cs -ls` In version 0.16.1, the `-a` option is added for output optimization, as shown in the following figure.

```
The pangu disk status:
Total Disk Size:           75044 GB
Total Free Disk Size:      73663 GB
Total Pangu Usable Disk Size: 75044 GB
Total Pangu Usable Free Size: 73663 GB
Total File Size:           0 GB
Total User Reserved Size:  0 GB
Total User Used Size:      0 GB
Total Garbage Size:        0 GB
Total Abnormal Size:       0 GB
Redundancy Ratio:         3
TotalChunkNumber:4167      NonTempChunkNumber:4167      NonTempChunkDataSize:0 GB
```

- Query abnormal chunks: `puadmin fs -abnchunk -t [none|onecopy|lessmin|lessmax]`
- Query which file the abnormal chunk belongs: `puadmin fs -whois 5973023903449089`
- Delete a file: `pu rm /systest/pangu /rs1d04281.et2sqa/RAFWriteRead/BlockSize_4096/16/0`
- Clear the recycle bin: `puadmin fs -crb -f`
- Query the election status: `puadmin gems`
- Check version consistency: `puadmin env-gbi-c`
- Query decommissioned machines: `./puadmin cs -ls --puadmin_ShowDecommissionChunkserver=true |grep tcp`

2.6.3.4 GC of Apsara Distributed File System

This topic describes how to enable and disable GC in puadmin and view the GC status of Apsara Distributed File System.

In Apsara Distributed File System, all deleted files are moved to the recycle bin first and removed from the recycle bin after a certain period (one day by default). The flag is `pangu_master_DelayTimeForFileGC`, in seconds. When the memory usage in the Apsara Distributed File System master exceeds the threshold (85% of the total memory), Apsara Distributed File System automatically clears the recycle bin to clear the deleted files.

To prevent deleted files from being permanently cleared from the disk, you can set `pangu_master_ForceFileGCThreshold` to 100 to disable this function. If the function is disabled, the following information is displayed: `Pangu master in mode : PANGU_MASTER_MODE_SAFE_OPEN & PANGU_MASTER_MODE_ALLOW_WRITE`.

You can implement the following GC functions of Apsara Distributed File System in puadmin:

- Enable GC

Run the following command to enable GC:

```
/apsara/deploy/puadmin ms -stat --safe=off
```

Command output:

```
Master Address: nuwa://localcluster/sys/pangu/master
```

```
Pangu master in mode: PANGU_MASTER_MODE_SAFE_CLOSE & PANGU_MASTER_MODE_ALLOW_WRITE
```



Note:

If the output contains `PANGU_MASTER_MODE_SAFE_CLOSE`, GC is enabled, which means that the files deleted from Apsara Distributed File System will be automatically recycled by GC.

- **Disable GC**

Run the following command to disable GC:

```
/apsara/deploy/puadmin ms -stat --safe=on
```

Command output:

```
Master Address: nuwa://localcluster/sys/pangu/master
```

```
Pangu master in mode: PANGU_MASTER_MODE_SAFE_OPEN & PANGU_MASTER_MODE_ALLOW_WRITE
```



Note:

If the output contains PANGU_MASTER_MODE_SAFE_OPEN, GC is disabled, which means that the deleted files on Apsara Distributed File System will not be automatically recycled by GC.

- **Check the GC status**

Run the following command to check the GC status:

```
/apsara/deploy/puadmin ms -stat -M
```

Command output:

```
Master Address: nuwa://localcluster/sys/pangu/master
```

```
Pangu master in mode: PANGU_MASTER_MODE_SAFE_CLOSE & PANGU_MASTER_MODE_ALLOW_WRITE
```



Note:

If the output contains PANGU_MASTER_MODE_SAFE_CLOSE, GC is enabled. If the output contains PANGU_MASTER_MODE_SAFE_OPEN, GC is disabled.

2.6.3.5 Cluster rebalance

Apsara Distributed File System implements even data storage during operation

. However, a newly added machine that is less occupied can cause data skew. In addition, the new machine carries heavy load. We recommend that you enable the rebalance function on the back end. You can run the following commands to enable the function in puadmin:

```
/apsara/deploy/puadmin rebalance -expand #Recompute data distribution.
```

```
/apsara/deploy/puadmin rebalance -start #Start rebalance.
```

```
/apsara/deploy/puadmin rebalance -stat #Check for data migration completed by rebalance.
```

```
/apsara/deploy/puadmin rebalance -stop #Stop rebalance.
```

**Note:**

Rebalance can only be stopped manually. We recommend that you stop rebalance when the disk usage of the machines in the cluster are close to each other after rebalance runs for a period.

The rebalance traffic is restricted to prevent the front-end read/write operations from being affected. To speed up the rebalance process, you can use the `pangu_chunkserver` flag to configure a larger traffic threshold (in Mbit/s). The configuration applies to the whole cluster. Run the following commands to increase or restore the traffic threshold:

Increase the traffic threshold:

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationMinReadNetThroughput 300 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationMaxReadNetThroughput 300 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationMinWriteNetThroughput 300 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationMaxWriteNetThroughput 300 -c
```

Restore the traffic threshold to the default value:

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationMinReadNetThroughput 10 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationMaxReadNetThroughput 30 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationMinWriteNetThroughput 10 -c
```

```
/apsara/deploy/puadmin flag -set pangu_chunkserver_ReplicationM  
axWriteNetThroughput 30 -c
```

2.6.3.6 Directory quota operations

Set quotas

Command

Run the following command to specify directory quotas in Apsara Distributed File System:

```
puadmin fs -quota <dir-name> --set=<entryCount,physicalSize,fileCount,  
logicalSize>
```

You can set four directory quotas: the number of subdirectories and files in a directory (entryCount), the physical size of a file (physicalSize), the number of files (fileCount), and the logical size of a file (logicalSize).

These quota values can be positive integers or strings "default" and "unlimited." The string "default" indicates that the quota of the entry remains unchanged. The string "unlimited" indicates that no quota is set for the entry.

Example

- ```
$puadmin fs -quota pangu://localcluster/testQuota/ --set=500,200,300,
default
```

**Set quotas for the testQuota directory. The maximum number of subdirectories and files in the directory is 500, the maximum physical size is 200, the maximum number of files is 300, and the maximum logical size remains unchanged.**

- ```
$puadmin fs -quota pangu://localcluster/testQuota/ --set=500,200,300,  
unlimited
```

Set quotas for the testQuota directory. The maximum number of subdirectories and files in a directory is 500, the maximum physical size is 200, the maximum number of files is 300, and the maximum logical size is unlimited.



Note:

- **Quota operations are directory specific. You need to add a forward slash (/) to the end of the directory to differentiate directories from files.**

- **The maximum number of files does not include replicas. However, the maximum physical file size is affected by replicas. For example, (3,3) indicates that the file size is three times of the original file size.**
- **If the system contains RaidFile, its logical size is accurate. The physical size of a single file may have a deviation up to several bytes due to rounding.**

View quotas

Command

```
puadmin fs -quota <dir name> or pu quota <dir name>
```

Run this command to query the directory quotas, including the specified quota under a directory and the actual number of directories, number of files, and file size.

Example

```
$/apsara/deploy/puadmin fs -quota pangu://localcluster/apsara/
```

Command output:

```
quota under pangu://localcluster/apsara/  
EntryNumber Limit:unlimited Used:3  
FileNumber Limit:unlimited Used:2  
FilePhysicalLength Limit:unlimited Used:626292888  
FileLogicalLength Limit:unlimited Used:2087642962
```

Delete quotas

Command

```
puadmin fs -quota <dir> -r
```

Run this command to delete the quotas of a directory. After this command is executed, the quotas of EntryNumber, FileNumber, FilePhysicalLength, and FileLogicalLength become unlimited.

Example

```
$/apsara/deploy/puadmin fs -quota pangu://localcluster/apsara/ -r
```

Command output:

```
quota under pangu://localcluster/apsara/  
EntryNumber Limit:unlimited Used:3  
FileNumber Limit:unlimited Used:2  
FilePhysicalLength Limit:unlimited Used:626292888
```

```
FileLogicalLength Limit:unlimited Used:2087642962
```

2.6.3.7 Directory pin operations

You can pin directories to prevent accidental deletion or renaming.



Note:

To pin a directory, you must pin its parent directory first. To unpin a directory, ensure that all its subdirectories are unpinned. In most cases, you can pin a high-level directory.

Pin a directory

- Use the following command to pin the full path of Apsara Distributed File System , which must end with a forward slash (/):

```
$puadmin fs -pin pangu://localcluster/apsara/
```

Command output:

```
Directory pinned!
```

- Run the following command to pin a local path of Apsara Distributed File System:

```
$puadmin fs -pin /apsara/deploy/
```

Command output:

```
No Master Address specify, using nuwa://localcluster/sys/pangu/
master
Directory pinned!
```



Note:

An error is returned when you try to delete a pinned directory. For example, if you run `$pu rmdir -f -p /apsara/`, error:Directory is pinned is returned.

Unpin a directory

Run the following command to unpin the directory containing the pin subdirectory :

```
$puadmin fs -unpin /apsara/
```

Command output:

```
No Master Address specify, using nuwa://localcluster/sys/pangu/master
--
unpin:Sub dirs are pinned, please unpin all first.
```

```
Pinned dirs:
deploy/
```

**Note:**

The following error is returned when you try to unpin the directory containing the pin subdirectory: Error: unpin:Sub dirs are pinned, please unpin all first
.

In this case, you must run the `$puadmin fs-unpin/apsara/deploy/` command to unpin the directory.

Command output:

```
No Master Address specify, using nuwa://localcluster/sys/pangu/master
Directory unpinned!
```

Check whether a directory is pinned

Run the `$pu dirmeta/apsara/` command to check whether the directory is pinned.

Command output:

```
pangu://localcluster//apsara/
Length      : 0
FileNumber  : 0
DirNumber   : 2
Pinned      : 1
```

**Note:**

If the value of Pinned is 1, the directory is pinned. If the value is 0, the directory is not pinned.

2.6.4 Operations and maintenance of masters

2.6.4.1 Overview

Apsara Distributed File System V0.16 supports multiple masters. O&M operations supports master switchover, status check, replacement, log synchronization, and volume expansion. At present, Apsara Distributed File System master supports Federation. All accesses are directed to DefaultVolume by default. To access a certain volume, specify the `-v` parameter.

For example, run the following command to list all the volumes under the master:

```
/apsara/deploy/puadmin vol -ls
```

Command output:

```
Name: v2  
Name: PanguDefaultVolume
```

This indicates that the master has two volumes. If the `-v` parameter is not specified, the command takes effect only on the PanguDefaultVolume volume. To operate on the v2 volume, specify the `-v v2` parameter. The commands are as follows:

```
/apsara/deploy/puadmin lscs: List the chunkserver of the default volume
```

```
/apsara/deploy/puadmin lscs-v v2: List the cs of the v2 volume
```

2.6.4.2 Switch over the primary master

Scenario

A master group consists of multiple masters. Only one master is in the primary state. When an exception occurs on the primary master, the PE assesses the situation and switches services from the primary master to a secondary master if necessary. Assume the current primary master is A, the secondary masters are B and C, and you need to switch services from A to B.

Command

```
/apsara/deploy/puadmin ms -sp B.tcpAddress
```

**Note:**

Run `/apsara/deploy/puadmin gems` to obtain the TCP address of B.

Operation results

- **If the operation is successful, `Switch primary form A.tcpAddress to B.tcpAddress succeed. is displayed.`**
- **If the operation fails, services are probably switched to C. You can run `/apsara/deploy/puadmin ms -elec` to check whether the primary master is switched over successfully.**

2.6.4.3 Status check for multiple masters

2.6.4.3.1 View the election status

Command

```
/apsara/deploy/puadmin ms -elec
```

Command output

```
ElectMasterStatus : ELECT_MASTER_OVER_ELECTION
PrimaryId         : tcp://10.101.164.1:10260
PreferedWorkerid  :
PrimaryLogId      : 882290350
TotalWokerNumber  : 3
ElectConsentNumber : 2
SyncConsentNumber : 2
ElectSequence     : [b2bca55e-530d-49ad-96d6-f3e564aece6d,1,
1782091301]
WorkerStatus      :
    tcp://10.101.164.10:10260 : ELECT_WORKER_STATUS_SECONDARY
    tcp://10.101.164.12:10260 : ELECT_WORKER_STATUS_SECONDARY
    tcp://10.101.164.1:10260  : ELECT_WORKER_STATUS_PRIMARY
```

Output description

The current environment has only one primary master and its TCP address is tcp://10.101.164.1:10260. (See tcp://10.101.164.1:10260 : ELECT_WORKER_STATUS_PRIMARY)

The current environment has two secondary masters and their TCP addresses are tcp://10.101.164.10:10260 and tcp://10.101.164.12:10260.

(See tcp://10.101.164.10:10260 : ELECT_WORKER_STATUS_SECONDARY and tcp://10.101.164.12:10260 : ELECT_WORKER_STATUS_SECONDARY)

2.6.4.3.2 View the log synchronization status of multiple masters

Command

```
/apsara/deploy/puadmin ms -elec -s
```

Command output

```
PrimaryStatus : PRIMARY_STARTUP_SERVICE_STARTED
PrimaryCurrentLogId : 882290350
WorkerSyncStatus :
tcp://10.101.164.10:10260[SyncedLogId:882290350, LastFailTime:1970-01-01 08:00:00, WorkerType: NORMAL, LogGap:0]
```

```
tcp://10.101.164.12:10260[SyncedLogId:882290350, LastFailTime:1970-01-01 08:00:00, WorkerType: NORMAL, LogGap:0]
```

Output description

The status log ID of the Pangu master is 882290350, and the log ID is a monotonically increasing sequence.

The secondary master can be in either the NORMAL or VIRTUAL state.

A master in the VIRTUAL state can only serve as a secondary master and receive logs synchronized from the primary master. It has no impact on the service.

2.6.4.4 Rules for the Apsara Distributed File System master to write .cpt and .log files

The .cpt and .log files recorded by Apsara Distributed File System master are named in the following formats: pangu_master_op.logId.cpt and pangu_master_op.logId.log. logId indicates the memory state when the .cpt file is created and the state of the log file switchover. Each .cpt file corresponds to a .log file with the same log ID. The .log file cannot be empty. Otherwise, the Apsara Distributed File System master cannot use the .cpt and .log file pair to recover data in memory.

For example, the following files exist:

/apsarapangu/pangu_master_op.673284360.cpt

/apsarapangu/pangu_master_op.673284360.log

/apsarapangu/pangu_master_op.673285365.log

/apsarapangu/pangu_master_op.673517452.cpt

/apsarapangu/pangu_master_op.673517452.log

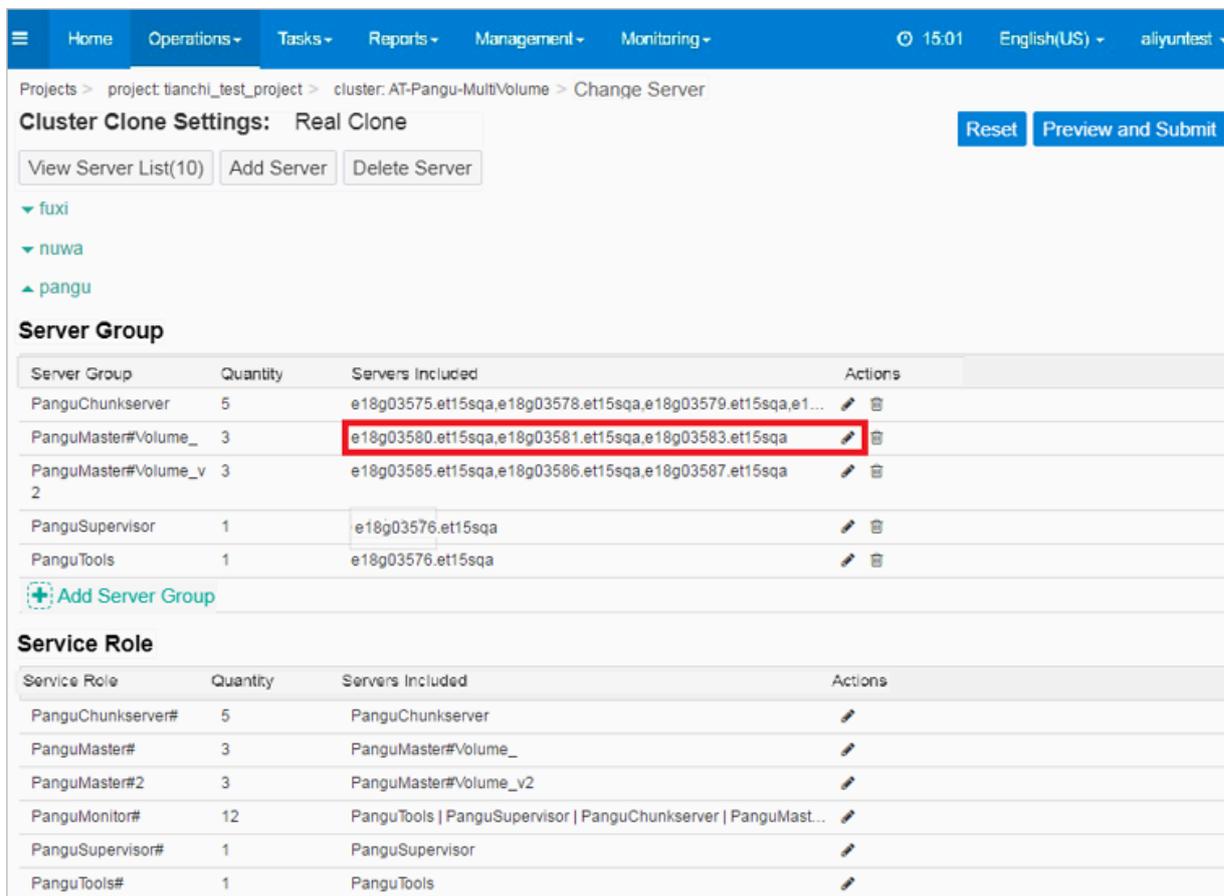
Then, pangu_master_op.673517452.cpt and pangu_master_op.673517452.log are the latest pair of .cpt and .log files.

2.6.4.5 Master replacement

2.6.4.5.1 Procedure

Apsara Distributed File System 0.15 and later versions support multiple volumes. Each master belongs to one volume.

To replace the master, choose **Operations > Cluster Operations** from the top navigation bar. Select the corresponding cluster and choose **Management > Machine Change** to go to the machine change page.



Projects > project: tianchi_test_project > cluster: AT-Pangu-MultiVolume > Change Server

Cluster Clone Settings: Real Clone Reset Preview and Submit

View Server List(10) Add Server Delete Server

▼ fluxi
▼ nuwa
▲ pangu

Server Group

Server Group	Quantity	Servers Included	Actions
PanguChunkserver	5	e18g03575.et15sqa,e18g03578.et15sqa,e18g03579.et15sqa,e1...	
PanguMaster#Volume_1	3	e18g03580.et15sqa,e18g03581.et15sqa,e18g03583.et15sqa	
PanguMaster#Volume_v2	3	e18g03585.et15sqa,e18g03586.et15sqa,e18g03587.et15sqa	
PanguSupervisor	1	e18g03576.et15sqa	
PanguTools	1	e18g03576.et15sqa	

+ Add Server Group

Service Role

Service Role	Quantity	Servers Included	Actions
PanguChunkserver#	5	PanguChunkserver	
PanguMaster#	3	PanguMaster#Volume_1	
PanguMaster#2	3	PanguMaster#Volume_v2	
PanguMonitor#	12	PanguTools PanguSupervisor PanguChunkserver PanguMast...	
PanguSupervisor#	1	PanguSupervisor	
PanguTools#	1	PanguTools	

The machines contained in PanguMaster# are specified by the machine group PanguMater#Volume_*. To replace the master, you only need to modify the machine group PanguMater#Volume_*. For example, the preceding figure has two volumes. To replace the master of the default volume, you only need to modify the PanguMaster#Volume_ machine group.



Note:

- The newly added machine must have no master oplogs in the `/apsarapangu` and `/apsarapangu/backup` directories, and its `/apsarapangu/conf/` directory must be empty.
- Only one master can be replaced at a time. After replacement, check that the Apsara Infrastructure Management Framework rolling task is completed using `puadmin` before starting the next replacement.

- **To ensure data security, do not replace three masters consecutively. If you need to replace three pangu masters consecutively, perform the following steps before each replacement:**
 1. **Run the `/apsara/deploy/putouch pangu://localcluster/xxx` command to generate one oplog.**
 2. **Send the `/apsara/deploy/puadmin ms -dump` command to three masters to enable each master to generate a checkpoint. Confirm that the new checkpoint has been generated on each master before performing the next operation. By default, the checkpoint of Apsara Distributed File System master is under the `/apsarapangu/` directory. The file name format is similar to `pangu_master_op.370748130.cpt`, but the numbers in the file name are different. Check that the file generation time is later than the command running time.**
 3. **Run the `/apsara/deploy/purmpangu://localcluster/xxx` command.**
 4. **Modify the configuration in the Apsara Infrastructure Management Framework console and perform replacement.**
- **If the master is damaged, you can replace it with a new one using the master replacement process.**
- **If a master is damaged, you can replace it. Only replacing damaged masters is allowed.**
- **Exchanging machines across groups is not allowed.**

For example, an environment has two volumes: `PanguMater#Volume_ : [1, 2, 3]` and `PanguMater#Volume_v2 : [4, 5, 6]`.

The server roles of the volumes are `PanguMaster#:` `PanguMater#Volume_ | PanguMater#Volume_v2`.

In this case, changing `[1, 2, 3] [4, 5, 6]` to `[1, 5, 3] [4, 2, 6]` is not allowed because it results in volume exceptions.

In some applications, the supervisor and the `pangu_master` are deployed on one machine. If you replace the machine and directly put it out of service, the supervisor may also become out of service. If all supervisors are out of service, the follow-up O&M operations are blocked because they cannot be approved by the supervisor. Therefore, you need to change the machine group of the supervisor during master replacement to prevent shutting down the supervisor.

**Note:**

The replacement of the `pangu_master` is a high-risk operation. Before one replacement is completed, do not start the next replacement. Check whether a replacement has been completed by referring to the following criteria:

- The rolling task of Apsara Infrastructure Management Framework is complete.
- In the Apsara Infrastructure Management Framework console, the number of active and out-of-service machines of the corresponding cluster is 0.
- Run the `/apsara/deploy/puamin gems` command on AG to confirm that the new machine is working normally and the out-of-service machine cannot be queried.

2.6.4.5.2 Replace a master

The master replacement process involves two steps, one for getting the new master into service and one for getting the existing master out of service. The supervisor is not informed when a new master is put into service. The supervisor only receives an application for getting a master out of service. The supervisor then sends an out-of-service command to the master. The primary master removes the existing master and returns OK to the supervisor only when a new master joins the cluster and logs are synchronized to the new cluster. After the supervisor receives the OK message, it sends an approval to Apsara Infrastructure Management Framework. Then, Apsara Infrastructure Management Framework brings the existing master out of service.

If the existing master is the primary master, the supervisor performs switchover before issuing the out-of-service command. If the new master is not in service, the out-of-service task for the existing master remains unapproved until the task times out.

This is because the supervisor approves the task after the new master goes into service and completes log synchronization. The time required for log synchronization depends on the size of the current `.cpt` file, the difference between the `.log` file and the `.cpt` file, and the load on the current master. At present, the bandwidth for `.cpt` file replication between masters is 50 Mbit/s. The `.cpt` file synchronization speed between masters is 35,000 entries per second. If the `.cpt` file size is 20 GB and the `.log` file contains 10 million entries, `.cpt` file synchronization for the new master takes $20 \times 1024/50 = 410$ seconds. If the load on the current master is 15,000

entries per second, it takes $1000/(3.5 - 1.5) = 500$ seconds to synchronize 10 million log entries. The total time is 910 seconds, that is, 15 minutes.

2.6.4.5.3 Manually replace a master

When automatic replacement fails, perform manual replacement after the new master starts properly.

Before manual replacement, verify that the supervisor is unable to perform approval with the self-service team. Assume that the IP address of the master to be brought into service is new, the IP address of the master to be put out of service is old, and the former has been started properly.

If it is not started properly, use Apsara Infrastructure Management Framework to deploy and start a new master. If the new master cannot be started, contact Apsara Infrastructure Management Framework support personnel.

After the new master is started properly, perform the following operations:

Perform the following steps on AG:

1. Run the `puadmin gss` command and check whether the new master is included in the command output, to determine whether the new master has been added to the master group. If the name of the new master is included in the command output, proceed with step 2. If it is not included in the command output, run the `/apsara/deploy/puadmin ms -elec --role --add=tcp://new:10260 -virtual yes` command to add the new master to the master group and set the state of the new master to virtual.
2. Run the `puadmin gss` command to view the synchronization status, as shown in the following figure.

```

$./puadmin gss
PrimaryStatus : PRIMARY STARTUP_SERVICE_STARTED
PrimaryCurrentLogId : 5068945
WorkerSyncStatus :
    tcp://100.81.240.140:10260 [SyncedLogId:5068945, LastFailTime:
    tcp://100.81.240.143:10260 [SyncedLogId:5068945, LastFailTime:

```

The red box in the figure above indicates the ID of the primary master, and the brown box indicates the log ID of the new machine. If the difference between the two is less than 10,000, log synchronization is completed.

3. After synchronization, run the `/apsara/deploy/puadmin ms -elec --role --add=tcp://old:10260 -virtual yes` command to set the state of the old master to virtual.
4. Run the `/apsara/deploy/puadmin ms -elec --role --add=tcp://new:10260 -virtual no` command to set the state of the new master to normal.
5. Run the `/apsara/deploy/puadmin ms -elec --role --rm=tcp://old:10260` command to delete the old master.

2.6.4.6 Manually synchronize logs between a primary master and a secondary master

Scenarios

Logs must be manually synchronized between a primary master and a secondary master when automatic log synchronization fails due to a large difference in log quantity resulting from a fault or incorrect operation.

Prerequisites

- This operation is performed without stopping the Apsara service.
- Multiple Apsara Distributed File System masters are running.

Procedure

1. Run the following command to generate the latest checkpoint file for the primary master:

```
/apsara/deploy/puadmin ms -dump
```

2. Copy the new `.cpt` and `.log` files generated in the primary master to the `/apsarapangu/` directory of the secondary master. Ensure that the latest `.cpt` and all the related logs are copied. Restart the `pangu` process on the secondary master.



Note:

The `/apsara/deploy/puadmin ms -dump` command sends commands to the primary master and returns whether the commands are sent successfully by default. Wait until the files are generated.

2.6.4.7 Rename a chunkserver online

Scenarios

The chunkserver name configured during cluster deployment on a machine where the chunkserver resides is incorrect. The chunkserver must be renamed without interrupting the service. In this case, you can use puadmin to change the chunkserver name recorded in the master memory.

For example, you can change the name of the chunkserver at tcp://10.101.164.7:10260 to rs1d04271.et2sqa_new.

Command for renaming the chunkserver

```
/apsara/deploy/puadmin ms -modify --csm -n tcp://10.101.164.7:10260,  
rs1d04271.et2sqa_new
```

Expected output

Modify chunkserver name success.

Command for querying the result

```
/apsara/deploy/puadmin cs -ls --service=tcp://10.101.164.1:10260 | grep  
tcp://10.101.164.7:10260
```

Expected output

```
31. NORMAL (80000/129228) (ttl= 20) tcp://10.101.164.7:10260 rs1d04271.  
et2sqa_new SendBuffer : 0(KB)
```



Note:

If there are multiple masters, confirm that the new chunkserver name meets expectations and the name has been changed on each master. After the change, a checkpoint is generated on each master.

Command for generating a checkpoint on a master

```
/apsara/deploy/puadmin ms -dump --Server=tcp://10.101.164.1:10260
```

Expected output

Start to generate checkpoint now.

2.6.4.8 Multi-master tools

Command	Description
<code>/apsara/deploy/puadmin ms -elec</code>	Views the election status.
<code>/apsara/deploy/puadmin ms -elec -m tcp://10.138.26.24:10260</code>	Views the ElectWorker status.
<code>/apsara/deploy/puadmin ms -elec -s</code>	Queries the status of Oplog synchronization among multiple masters.
<code>/apsara/deploy/puadmin ms -sp tcp.Address</code>	Switches over a primary master.
<code>/apsara/deploy/puadmin ms -elec -M tcpAddress</code>	Displays all the ElectWorker metrics and their values of Apsara Distributed File System.
<code>/apsara/deploy/puadmin ms -elec -l</code>	Obtains the latest Oplog.

2.6.5 Operations and maintenance of chunkservers

2.6.5.1 Set the chunkserver status

You can manually set the chunkserver status. The available states are **NORMAL**, **READONLY**, and **SHUTDOWN**. The following example shows how to set the chunkserver state to **SHUTDOWN**.

Command

```
/apsara/deploy/puadmin cs -stat tcp://10.101.164.7:10260 --set=SHUTDOWN
```

Expected output

```
set chunkserver status: SHUTDOWN
```

Query the setting result

Run the `/apsara/deploy/puadmin lscs | grep tcp://10.101.164.7` command to query the setting result.



Note:

The query takes effect only after the master detects that the chunkserver enters the **SHUTDOWN** state. This takes less than 1 minute.

Expected output

```
31. SHUTDOWN (NA) (ttl= 20) tcp://10.101.164.7:10260 rs1d04271.
et2sqa_new SendBuffer : 0(KB), Backup: Doing
```

2.6.5.2 Set the disk status

You can manually set the chunkserver disk status. The available states are OK, SHUTDOWN, and ERROR. The following example shows how to set the disk state to DISK_ERROR.

Command

```
/apsara/deploy/puadmin cs -stat tcp://10.101.164.7:10260 -d 1 --set=
ERROR
```

Expected output

```
set disk status success. After set : DISK_ERROR
```

Query the setting result

Run the `/apsara/deploy/puadmin cs -stat tcp://10.101.164.7:10260 -d 1` **command to query the setting result.**

Expected output

```
DiskId:1 DiskStatus:DISK_ERROR
```



Note:

If a cluster involves both the storage and journal scenarios, do not change the status of the disks and SSDCache from ERROR to OK. Instead, use the function that automatically brings empty disks into service, which is enabled by default.

If the status of a disk is changed from ERROR to OK in the preceding scenario, the following output is generated:

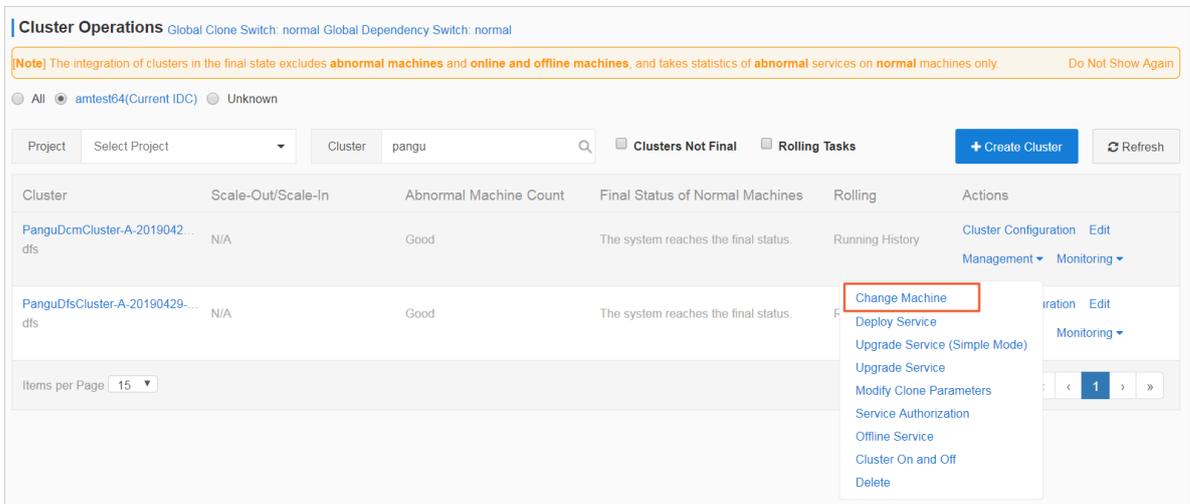
```
[admin@e18g06550 /apsarapangu/disk5/zhousu/dailycores/core_10.2017011117]
$/apsara/deploy/puadmin cs -stat tcp://100.81.240.125:10260 -d 1 --set=ok
Set disk status failed. Disk status: DISK_ERROR
--
Error: stat: For disk status, only OK/SHUTDOWN/ERROR are supported. READONLY is not supported now.
For chunkserver status, only NORMAL/READONLY/SHUTDOWN are supported.
```

2.6.5.3 Scale-out and scale-in of chunkservers

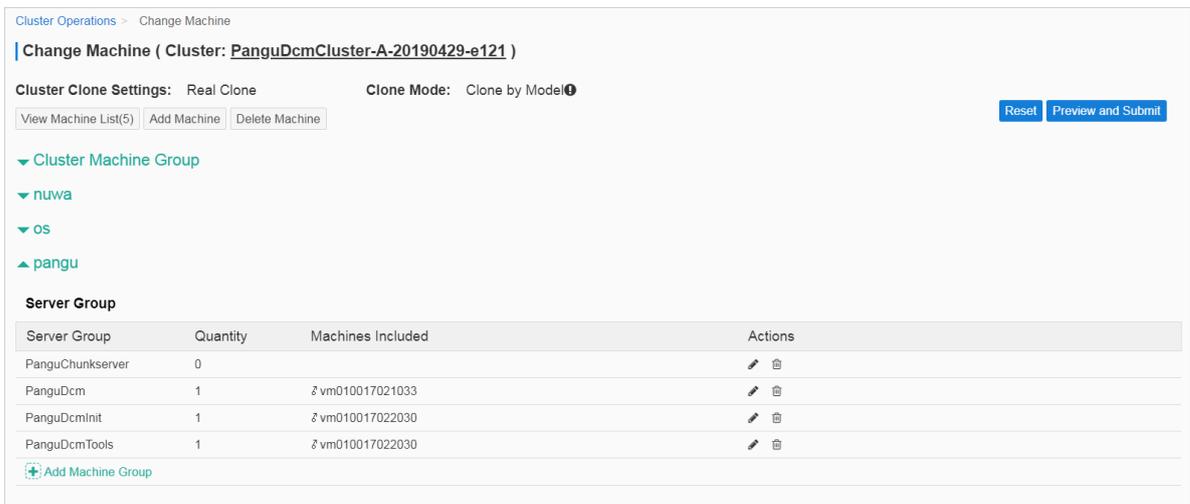
2.6.5.3.1 Procedure

This topic describes how to scale a chunkserver by using machine change.

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose Products > Apsara Infrastructure Management Framework.
3. Choose Operations > Cluster Operations to go to the cluster operations page.
4. Find the target cluster and choose Management > Machine Change to go to the machine change page, as shown in the following figure.



5. Change the list of machines that PanguChunkserver# points to. Add machines to the list to complete scale-out. Remove machines from the list to complete scale-in. Then, submit your changes, as shown in the following figure.



Note:

The added machines must belong to the current cluster. The machines in the default cluster can be used by the current cluster only after they are added to the current cluster by scaling out the cluster.

2.6.5.3.2 Scale-out procedure

The supervisor does not need to participate in chunkserver scale-out. You can scale out the chunkserver by adding machines to the machine list of the chunkserver server role. The machines must belong to the target cluster. If not, add them to the cluster first by scaling out the cluster. Apsara Infrastructure Management Framework deploys and starts the chunkserver on the machine.



Note:

- **The started chunkserver must be registered with Apsara Name Service and Distributed Lock Synchronization System. The NuwaConfig server role and NuwaLib server role must be deployed on the machine before scale-out, to ensure correct Apsara Name Service and Distributed Lock Synchronization System configuration. If the added machine was previously used by another cluster and still contains the data of that cluster, and do not have the NuwaConfig and NuwaLib server roles deployed, the chunkserver may read the existing data and register with Apsara Name Service and Distributed Lock Synchronization System of another cluster.**
- **Make sure that the added machine does not contain /apsarapangu/pangu_chunkserver_op*, which contains the metadata of the previous chunkserver. In the test environment, a new chunkserver may contain metadata. As a result, the chunkserver enters its original state after being started. In this case, clear the metadata before scaling out the chunkserver on that machine.**
- **If there are multiple zones, add the new chunkserver to the zones after scale-out. Otherwise, the chunkserver cannot be used.**

2.6.5.3.3 Scale-in procedure

The procedure of scaling out chunkservers is similar to that of scaling out chunkservers. You need to the modify the machine list of the chunkserver server role. To complete scale-in, you just need to remove chunkservers from the list as needed.

The scale-in operations in Apsara Infrastructure Management Framework must be approved by the supervisor. After receiving a chunkserver disconnection task, the supervisor sends the SHUTDOWN command to the corresponding chunkserver, waits for chunkserver data replication to complete, and then sends the approve command to Apsara Infrastructure Management Framework. After receiving the approve command, Apsara Infrastructure Management Framework stops the chunkserver process. The supervisor keeps sending the decommission command to the master to delete the chunkserver from the cluster. If the chunkserver process stops and the chunkserver enters the DISCONNECTED state, the decommission command is executed. The chunkserver is removed from the cluster.

The number of chunkservers that can be removed at a time depends on a flag of the supervisor. The flag is named `pangu_supervisor_MaxConcurrentChunkserverShutdownCount` and has the default value 1, indicating that one chunkserver can be removed at a time. To remove multiple chunkservers at a time, change the value.

The period from the time when a chunkserver process is stopped to the time when the chunkserver enters the DISCONNECTED state depends on two flags of `pangu_master`, `pangu_chunkserver_normal_ttl` and `pangu_chunkserver_disconnecting_ttl`. The default value of the first flag is 4, while the default value of the second is 16, so the total period is $(4 + 16) \times 15/60 = 5$ minutes. Set the first flag on `MaxCompute` to 40 so that the total period changes to 14 minutes. This means that a chunkserver is completely out of service at least 5 minutes or 14 minutes after data replication.

Duration of completing chunkserver disconnection

The supervisor performs approval only after chunkserver data replication is complete. The data replication time depends on the number of nodes in a cluster and the amount of data on the chunkserver to be disconnected. Currently, the maximum data replication speed is limited to 30 Mbit/s. Taking AY44B, a typical online OSS instance as an example, the capacity of most machines is close to full capacity, the disk capacity is 44 TB, and the cluster has 500 chunkservers. The maximum replication bandwidth is 30 Mbit/s. When all machines are involved in the replication, the estimated time for a single machine is $44 \times 1024 \times 1024 / (30 \times 500) = 3075$ seconds, about 50 minutes. If the machine has a high network load and cannot perform the replication at the maximum bandwidth, replicating data at the minimum replication bandwidth may take up to 150 minutes. If most of the data are

RAID files (8 + 3), the time is multiplied by 8 and increased to 400 to 1200 minutes. Therefore, the estimated time of Apsara Distributed File System data replication for disconnecting a chunkserver under AY44B is 50 to 150 minutes without RAID files, or 400 to 1200 minutes with RAID files (8 + 3). If the current state of the chunkserver is DISCONNECTED, the supervisor directly approves the request.

**Note:**

Before disconnecting a chunkserver, add the chunkserver to the blacklist of Job Scheduler. Otherwise, when TempFile is used, Job Scheduler may distribute instances to chunkservers in the SHUTDOWN state, which results in a failure to write TempFile.

When a chunkserver is disconnected, the supervisor automatically adds the chunkserver to the blacklist of Apsara Distributed File System, so that new chunks will no longer be distributed to the chunkserver. You can run the following commands to set and view the blacklist:

- **Set the blacklist:** `/apsara/deploy/puadmin upgrade -util -sbcs "tcp://x.x.x.x:10260, tcp://x.x.x.y:10260"`
- **View the blacklist:** `/apsara/deploy/puadmin upgrade -util -gbcs`

2.6.5.3.4 Disable manual scale-out of chunkservers by using puadmin

To scale in chunkservers in Apsara Infrastructure Management Framework, perform the standard scale-in procedure instead of manually running puadmin to shut down chunkservers. After data replication is complete, run the puadmin command to decommission the chunkserver from the cluster. Otherwise, the chunkserver list in Apsara Infrastructure Management Framework is inconsistent with the chunkserver list saved in pangu_master. For example, the number of chunkservers recorded in pangu_master is 100 while 120 in the Apsara Infrastructure Management Framework. When you operate on 20 chunkservers that do not exist in pangu_master, the supervisor will determine these chunkservers are not in the cluster and reject approval for security reasons.

Therefore, do not manually decommission chunkservers by using puadmin. Use the standard scale-in procedure instead.

2.6.6 Track the cluster status

2.6.6.1 View the cluster status

After initiating a task in Apsara Infrastructure Management Framework, you can view the status of the task in the Apsara Infrastructure Management Framework console.

1. Log on to the Apsara Stack Operations console.
2. In the left-side navigation pane, choose **Products > Apsara Infrastructure Management Framework** to go to the homepage of Apsara Infrastructure Management Framework.
3. Click **Homepage** to view the list of tasks that are in the normal, rollback, and paused states.

Click a state to view the cluster name, server role, and Git version of the task.



2.6.6.2 Submission history

You can view the history of operations on a cluster in Apsara Infrastructure Management Framework. Choose **Operations > Cluster Operations** from the top navigation bar. Select the corresponding cluster and choose **Monitoring > Operation Logs**, as shown in the following figure.

Cluster Operations Global Clone Switch: normal Global Dependency Switch: normal

[Note] The integration of clusters and final service status excludes abnormal servers and strong-online and offline servers, and takes statistics of abnormal services on normal servers only. Do Not Show Again

All amtest61(Current IDC) Unknown

Project: Select a project. Cluster: Select a cluster. Clusters Not Final Rolling Tasks + Create Cluster Refresh

Cluster	Scale-Out/Scale-In	Abnormal Server Count	Final Status of Normal Servers	Rolling	Actions
AcsControlCluster-A-2019042... acs	N/A	Good	The system reaches the final status.	Running History	Cluster Configuration Edit Management Monitoring
AlgoMarketCluster-A-2019051... pai	N/A	Good	The system reaches the final status.	Runr	Cluster Operation and Maintenance Center Service Final Status Query Operation Logs Distributed System Performance Monitor

On the page that appears, you can perform the following operations:

- View the cluster operation history, as shown in the following figure.

Cluster Operation Log (Cluster: AcsControlCluster-A-20190423-6fbf)					Refresh
Git Version	Description	Submitter Submitted At	Task Status	Actions	
b9f1af06300ed1b24a49035c9e52fb8b1d7799dc	auto update buildid.	aliyuntest 05/16/19, 18:27:09	Upgrade	View Release Changes	
ed7330560b7dea13b1d987a851698d33af4e307a	auto update buildid.	aliyuntest 05/15/19, 14:05:36	Upgrade	View Release Changes	

- Select a version number and click View Version Differences to compare the differences between any two submissions, as shown in the following figure.

Version Difference

Cluster Name: AcsControlCluster-A-201... Version: b9f1af06300ed1b24a49035c9e... Submitter: aliyuntest Description: auto update buildid

Select Base Version: Configuration Type Extend Configuration Cluster Configuration

Differential File List Total Differential Files: 3 (A: Add M: Modify R: Remove)

- M services/acs-accs_control/version.conf
- M services/tianji-sshtunnel-client/version.conf
- M tag.conf

- Select a base version and click Obtain Differences to view the difference details and track the content of each change, as shown in the following figure.

Cluster Name: AcsControlCluster-A-201... Version: b9f1af06300ed1b24a49035c9e... Submitter: aliyuntest Description: auto update buildid.

Select Base Version: Configuration Type Extend Configuration Cluster Configuration

Current File: services/acs-accs_control/version.conf

Base Version: 63e01abd0d70fc025762ee8f4168ec76aba215f9 Current Version: b9f1af06300ed1b24a49035c9e52fb8b1d7799dc

services/acs-accs_control/version.conf CHANGED

<pre> @@ -18,9 +18,9 @@ 18 Applications: 19 *: pangu#1184889_acs-accs_control_ControlInit_531c935c 20 CosConsoleAliyunCom#: 21 Applications: 22 - *: pangu#1187436_acs-accs_control_CosConsoleAliyunCom_531c935c 23 CrAuthentication#: 24 Applications: 25 *: pangu#1135306_acs-accs_control_CrAuthentication_531c935c 26 CrConsole#: @@ -30,9 +30,9 @@ 30 Applications: 31 *: pangu#1806050932381421 32 CrManager#: 33 Applications: 34 - *: pangu#1180197_acs-accs_control_CrManager_531c935c 35 CrOpenapi#: 36 Applications: </pre>	<pre> 18 Applications: 19 *: pangu#1184889_acs-accs_control_ControlInit_531c935c 20 CosConsoleAliyunCom#: 21 Applications: 22 + *: pangu#1188992_acs-accs_control_CosConsoleAliyunCom_531c935c 23 CrAuthentication#: 24 Applications: 25 *: pangu#1135306_acs-accs_control_CrAuthentication_531c935c 26 CrConsole#: @@ -30,9 +30,9 @@ 30 Applications: 31 *: pangu#1806050932381421 32 CrManager#: 33 Applications: 34 + *: pangu#1187940_acs-accs_control_CrManager_531c935c 35 CrOpenapi#: 36 Applications: </pre>
--	--

2.6.7 FAQ

2.6.7.1 Identify the machine running pangu_supervisor and the log location

On AG, run the `/apsara/deploy/nuwa_console --address=nuwa://localcluster/sys/pangu/supervisor` command to obtain the IP address of the machine where the supervisor runs.

The log path of pangu_supervisor is `/apsara/pangu_supervisor/log/pangu_supervisor.LOG`

2.6.7.2 Adjust the flag of pangu_supervisor

If you need to adjust the flag of the supervisor in the Apsara Infrastructure Management Framework console during O&M, wait until the current rolling ends. If you need the change to take effect immediately, set the flag by using the `puadmin` command. The flag set by using the `puadmin` becomes invalid upon restart.

Note that `puadmin` automatically increases the port number by 1 due to historical reasons. If the default listener port number of the supervisor is 10263, change it to 10262 when using the `puadmin` command to adjust the flag of the supervisor.

To view the value of a supervisor flag, run the following command: `/apsara/deploy/puadmin flag -get pangu_supervisor_MaxTolerantFailedChunkserver -s tcp://supervisor_ip_addr:10262`

To set the flag value, run the following command: `/apsara/deploy/puadmin flag -set pangu_supervisor_MaxTolerantFailedChunkserver 100 -s tcp://supervisor_ip_addr:10262`

2.6.7.3 Supervisor approval errors

2.6.7.3.1 Supervisor approval prerequisites

Before obtaining approval from the supervisor, the following conditions must be met: the cluster and master group are normal, all the chunkservers are in the NORMAL state, the cluster does not contain `abnchunks.`, and inter-cluster data replication is not in process. The normal causes and solutions are as follows:

- `CLUSTER_TYPE` in `tag.conf` is wrong, causing the supervisor fails to start. If the supervisor fails to start, you can run the `/apsara/pangu_supervisor/start` command and check whether any noticeable error is contained in the command output.
- The master group has no primary master. Check whether Apsara Name Service and Distributed Lock Synchronization System is available.
- The master logs are not synchronized and the number of logs of the primary master may differ greatly from that of the secondary master. The default number is 100.

- A `.cpt` file has not been created for masters for a long time and the `.cpt` ID may differ greatly from the log ID. The default value is 1048576.
- The number of abnormal chunkservers in the cluster exceeds the limit (10 for MaxCompute and 2 for others) specified by `pangu_supervisor_MaxTolerantFailedChunkserver`. This value affects hot upgrades, but does not affect chunkserver disconnection. Before an upgrade, add the abnormal chunkservers to the blacklist of the supervisor. The chunkservers in the blacklist are not included in the statistics on abnormal chunkservers. To add them to the blacklist, run the following command: `/apsara/deploy/puadmin flag -set pangu_supervisor_ChunkserverBlackList "csip1,csip2" -s tcp://supervisor_ip_addr:10262`.

**Note:**

Separate multiple chunkserver IP addresses with commas (,). Do not include tcp or port numbers.

- The cluster contains abnchunks. You can run the `/apsara/deploy/puadmin fs -abnchunk -t lessmin` command to query abnchunks. The logs of Pangu_supervisor contain "HasAbnormalChunkserverForMasterVolume."
- After the supervisor approves the masters during migration, masters in V0153 is terminated by Apsara Infrastructure Management Framework and it takes up to half an hour for masters in V016 to start. Although the master process exists, it stays in the DISCONNECTED state for a long time. This problem occurs when the disk speed is low and `pangu_OperationLogSyncDisk` is set to true. Before you perform upgrades, set this value to false in the template to accelerate startup.
- The number of replica tasks of the cluster exceeds the limit specified by `pangu_supervisor_DefaultOngoingReplicaForUpgrade`, and approval congestion occurs. To query the tasks, run the following command: `/apsara/deploy/puadmin rep -stat`.
- On some machines, a hostname may vary across volumes, while the corresponding IP address does not. In this case, you can run the `puadmin lscs -`

```
v volumename | grep tcp | awk '{print $6 $7}' | sort
```

command to check whether a hostname varies across volumes.

If yes, perform the following steps:

1. Run the `curl "127.0.0.1:7070/api/v3/column/m.ip? m.ip=10.101.162.176"` **command to view the hostname of the IP address in Apsara Infrastructure Management Framework.**
 2. Refer to [Rename a chunkserver online](#) and change the hostnames for consistency across volumes.
- Some server roles of `pangu_master` are not in the GOOD state. To prevent the dual masters problem after a task is stopped and then restarted, the supervisor checks the version and status of the `pangu_master` server role before approval. If it is not in the GOOD state, the supervisor rejects the request. In most cases, `pangu_monitor` generates an Error alert, making the server role enter the PROBATION state. The supervisor has a flag that can be used to skip checking the server role status. You can use the flag if you are certain about the server role status. To use the flag, set `pangu_supervisor_EnableCheckServerRoleState` to `false`.

2.6.7.3.2 Failure to approve the master replacement

The possible causes are as follows:

- The new master is not started.
- The logs of the new master have not been synchronized.

2.6.7.3.3 Failure to approve chunkserver disconnection

If the prerequisites have been met, the possible cause is that the number of chunkservers is smaller than the `mincopy` value and data replication fails. You can follow the following steps to resolve the issue:

1. **Identify the files that have not been copied:** If a chunkserver fails to disconnect for a long time, access `/apsara/pangu_chunkserver/log` on the chunkserver that is in the SHUTDOWN doing state, open `pangu_chunkserver.LOG`, and check whether the value of the `safeRemove` field is true or false. If it is false, search the previous log for the ID of the file related to the disconnection failure.

2. **Check the mincopy value: On AG, run the `/apsara/deploy/puadmin gfi FileId` command and search for the minCopy value in the file. If the value is greater than the current number of chunkservers, data cannot be replicated.**
3. **Run the `/apsara/deploy/puadmin whois FileId` command to convert the file ID into a file name.**
4. **Run the `/apsara/deploy/pu setreplica filename 3 5` command to set mincopy and maxcopy to 3 and 5, respectively.**

2.6.7.4 Accelerate chunkserver disconnection

Chunkservers are disconnected in series by default. To accelerate disconnection, you can change the value of `pangu_supervisor_MaxConcurrentChunkserverShutdownCount`. For details, see [Adjust the flag of pangu_supervisor](#). The value indicates the number of chunkservers that can be disconnected at the same time. If it is set to 100, a maximum of 100 chunkservers can be disconnected at the same time.

2.6.7.5 Rolling failure during a hot upgrade of Apsara Distributed File System

If you initiate a hot upgrade but the final rolling fails despite the supervisor's approval, the most possible cause is that an ERROR alert is generated by `pangu_monitor` during probation. The Apsara Infrastructure Management Framework console provides a list of machines reporting errors. You can search `/apsara/pangu_monitor/monitor/log/monitor.LOG` on the corresponding machines by the level.*error keyword.

2.6.7.6 Manually replace binary files in an emergency

2.6.7.6.1 Manual overwrite operation

Apsara Infrastructure Management Framework protects its files and automatically fixes changed files. This will affect the manually replaced binary files. Apsara Infrastructure Management Framework provides an overwrite tool to address this problem. Take `pangu_supervisor` as an example. The procedure is as follows:

1. **Copy the deployment directory of `pangu_supervisor`. Assume that `buildid` is 2861, the command is as follows:**

```
cp /cloud/app/pangu/PanguSupervisor#/pangu_supervisor/2861  
/cloud/app/pangu/PanguSupervisor#/pangu_supervisor/overwrite -rf
```

2. Create a configuration file under `service_manage`. The configuration file is named in the `Service name.Server role name.Application name` format. An example of a full path is as follows: `/cloud/data/tianji/TianjiClient#/service_manager/overwrite.d/pangu.PanguSupervisor#.pangu_supervisor`.

The file content is as follows:

```
{
  "service_name": "pangu",
  "sr_name": "PanguSupervisor#",
  "app_name": "pangu_supervisor",
  "against_work_dir": "/cloud/app/pangu/PanguSupervisor#/pangu_supervisor/2861",
  "work_dir": "/cloud/app/pangu/PanguSupervisor#/pangu_supervisor/overwrite",
  "expired_time": 9456814356
}
```

The parameters are described as follows:

- `against_work_dir`: the normal working directory.
 - `work_dir`: the temporary working directory used for replacement. The files in this directory are not affected by the downloader of Apsara Infrastructure Management Framework and are not automatically fixed by Apsara Infrastructure Management Framework.
 - `expired_time`: the timeout period. Set this parameter correctly so that the overwrite operation can take effect. This value is a UNIX timestamp.
3. After the file is saved, the SM restarts the supervisor. If you want to overwrite files on masters or chunkservers, perform the operation on the masters or chunkservers individually or in batches. If you perform the overwrite operation on them in batches, some processes are restarted at the same time, which may interrupt services. `apsara/pangu_supervisor` points to the overwritten directory. Apsara Infrastructure Management Framework does not fix the files changed in this directory. You can directly replace with the target binary files.
 4. After overwriting, the updated configuration is pushed, and the SM starts processes normally. After a version upgrade, the overwrite operation expires.

2.6.7.6.2 Use the overwrite tool of Apsara Infrastructure Management Framework

Apsara Infrastructure Management Framework provides an overwrite tool to simplify the overwrite procedure. The following example shows how to use this tool.

The version to be overwritten is 2899 and the validity period is 9,999 days. Perform the following steps:

1. Copy the directory.
2. Replace the binary files in the copied directory.
3. Perform the overwrite operation.

Take `pangu_chunkserver` as an example. The procedure is as follows:

1. Run the `ls -l` command in the `/apsara` directory to confirm that the chunkserver deployment directory is `/cloud/app/pangu/PanguChunkserver#/pangu_chunkserver/2899`. Run the following command to copy the directory: `cp /cloud/app/pangu/PanguChunkserver#/pangu_chunkserver/2899 /cloud/app/pangu/PanguChunkserver#/pangu_chunkserver/2899.overwrite -rf`
2. Replace the binary files.
3. Run the following command to overwrite the old version: `/cloud/tool/tianji/overwrite add pangu PanguChunkserver# pangu_chunkserver 2899 9999`



Note:

The chunkserver is restarted during this step and switched to the new version.

4. In the `/apsara` directory, run the `ls -l` command to check whether overwriting is successful. If the output ends with `.overwrite`, overwriting is successful.
5. Run the following command to manually cancel overwriting: `/cloud/tool/tianji/overwrite remove pangu PanguChunkserver# pangu_chunkserver`. After the command is executed, the overwrite operation does not take effect.

2.7 ApsaraDB for RDS

2.7.1 Architecture

2.7.1.1 System architecture

2.7.1.1.1 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, making the data more traceable.

Automatic backup

RDS provides various types of backup. MySQL instances support both physical and logical backup. PostgreSQL and PPAS instances support full backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can 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.

ApsaraDB RDS for PostgreSQL and ApsaraDB RDS for PPAS automatically perform full physical backup.

Data tracing

ApsaraDB for RDS can use the backup files and logs to generate a temporary instance from any time point within seven days. Verify that your data does not have any errors before you restore it.

Creating a temporary instance will not affect the operations of the current instance.

You can create only one temporary instance at a time for each RDS instance. A maximum of 10 temporary instances can be created each day. Each temporary instance is valid for 48 hours.

2.7.1.1.2 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.

2.7.1.1.3 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.

2.7.1.1.4 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.

The HA service uses the Detection, Repair, and Notice modules to ensure the availability of data link services. It also processes internal database exceptions.

HA policies

Each HA policy defines a combination of service priorities and data replication modes defined to meet the needs of your business.

There are two service priorities:

- **Recovery time objective (RTO):** The database preferentially restores services to maximize the availability time. Use the RTO policy if you require longer database uptime.
- **Recovery point objective (RPO):** The database preferentially ensures data reliability to minimize data loss. Use the RPO policy if you require high data consistency.

There are three data replication modes:

- **Asynchronous replication (Async):** When an application initiates an update request such as add, delete, or modify operations, the primary node responds to the application immediately after the primary node completes the operation. The primary node then replicates data to the secondary node asynchronously. This means that the operation of the primary database is not affected if the secondary node is unavailable. Data inconsistencies may occur if the primary node is unavailable.
- **Forced synchronous replication (Sync):** When an application initiates an update request such as add, delete, or modify operations, the primary node replicates data to the secondary node immediately after the primary node completes the operation. The primary node then waits for the secondary node to return a success message before the primary node responds to the application. The primary node replicates data to the secondary node synchronously. Unavailability of the secondary node will affect the operation on the primary node. Data will remain consistent even when the primary node is unavailable.

- **Semi-synchronous replication (Semi-Sync):** Data is typically replicated in Sync mode. When trying to replicate data to the secondary node, if an exception occurs causing the primary and secondary nodes to be unable to communicate with each other, the primary node will suspend response to the application. If the connection cannot be restored, the primary node will degrade to Async mode and restore response to the application after the Sync replication times out. In a situation such as this, the primary node becoming unavailable will lead to data inconsistency. After the secondary node or network connection is recovered, data replication between the two nodes is resumed, and the data replication mode will change from Async to Sync.

You can select different combinations of service priorities and data replication modes to improve availability based on the business features.

2.7.1.1.5 Task scheduling system

You can use the RDS console or APIs to create and delete instances, or switch instances between the internal and public networks. All instance operations are scheduled, traced, and displayed as tasks.

Resource

The Resource module allocates and integrates lower-level RDS resources to enable and migrate instances. For example, when you use the RDS console or an API to create an instance, the Resource module calculates which physical server is most suitable to bear traffic. This module also allocates and integrates lower-level resources required to migrate RDS instances. After instances are created, deleted, and migrated, the Resource module calculates the fragmentation of resources, and periodically integrates resource fragments to handle traffic spikes.

2.7.2 RDS O&M overview

Apsara Stack Operations Console provides the following RDS O&M features:

- **Instance management:** allows you to view the details, logs, and user information of an instance.
- **Host management:** allows you to view and manage hosts.

2.7.3 Log on to Apsara Stack Operations console

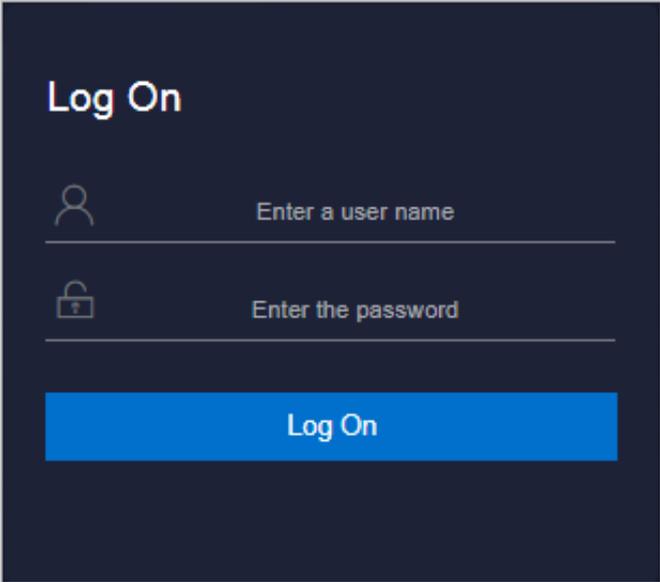
Prerequisites

- 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.

Figure 2-10: 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.

2.7.4 Manage instances

You can view instance details, logs, and user information.

Procedure

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

3. On the Instance Management tab of RDS, you can view the following information:

- **Instances**

On the Instance Management tab, view the instances under the account, as shown in *Figure 2-11: Instances*.

Figure 2-11: Instances

Instance Name	Availa...	CPU Perfor...	QPS Perfor...	IOPS Perfor...	Conne...	Disk Usage	Instance Status	Datab... Type	Actions
...	Yes				0		Creating	mysql	User Information Create Backup
...	Yes		2 %		0		Using	redis	User Information Create Backup

- **Instance details**

Click the ID of an instance to view its details, as shown in *Figure 2-12: Instance details*.

Figure 2-12: Instance details

Instance Information

Instance Name: m-...	CPU Performance: 0 %
Active-Standby Delay: 0	QPS Performance: %
Connections: 0	IOPS Performance: 0 %
Traffic:	Active Threads: 0
Client Instance Level: P4	Instance Status: █
Database Version: 5.6	Link Type: lvs
Cluster: ...	Created At: 09/27/2019, 16:12:54

Network Details of Instance Host

Host IP Addresses: ...	Proxies:
VIP IB_ID List of SLB: ...	ECS-typed Dedicated Host of Client Instance: No

Network Details of Instance-Attached Host

Host IP Addresses: ...	Proxies:
VIP IB_ID List of SLB: ...	ECS-typed Dedicated Host of Client Instance: No

Primary/Secondary Switch Query History

- **User information**

Click **User Information** in the **Actions** column, as shown in *Figure 2-13: User information*.

Figure 2-13: User information

The screenshot shows a 'User Information' page with a table of instance details. The table has the following columns: Instance Name, Instance Status, Database Type, Instance Usage Type, CPU Utilization, IOPS Utilization, Disk Utilization, and Connections Utilization. All instances listed are in the 'CREATING' status and use 'Redis' as the database type. Each row includes progress bars for usage and utilization metrics.

Instance Name	Instance Status	Database Type	Instance Usage Type	CPU Utilization	IOPS Utilization	Disk Utilization	Connections Utilization
[Redacted]	CREATING	Redis	[Redacted]	[Progress Bar]	[Progress Bar] %	[Progress Bar] %	[Progress Bar] %
[Redacted]	CREATING	Redis	[Redacted]	[Progress Bar]	[Progress Bar] %	[Progress Bar] %	[Progress Bar] %
[Redacted]	CREATING	Redis	[Redacted]	[Progress Bar]	[Progress Bar] %	[Progress Bar] %	[Progress Bar] %
[Redacted]	CREATING	Redis	[Redacted]	[Progress Bar]	[Progress Bar] %	[Progress Bar] %	[Progress Bar] %
[Redacted]	CREATING	Redis	[Redacted]	[Progress Bar]	[Progress Bar] %	[Progress Bar] %	[Progress Bar] %
[Redacted]	CREATING	Redis	[Redacted]	[Progress Bar]	[Progress Bar] %	[Progress Bar] %	[Progress Bar] %
[Redacted]	CREATING	Redis	[Redacted]	[Progress Bar]	[Progress Bar] %	[Progress Bar] %	[Progress Bar] %

4. On the **Instance Management** tab, click **Show Statistics** to view instance information by version and region, as shown in the following figure.



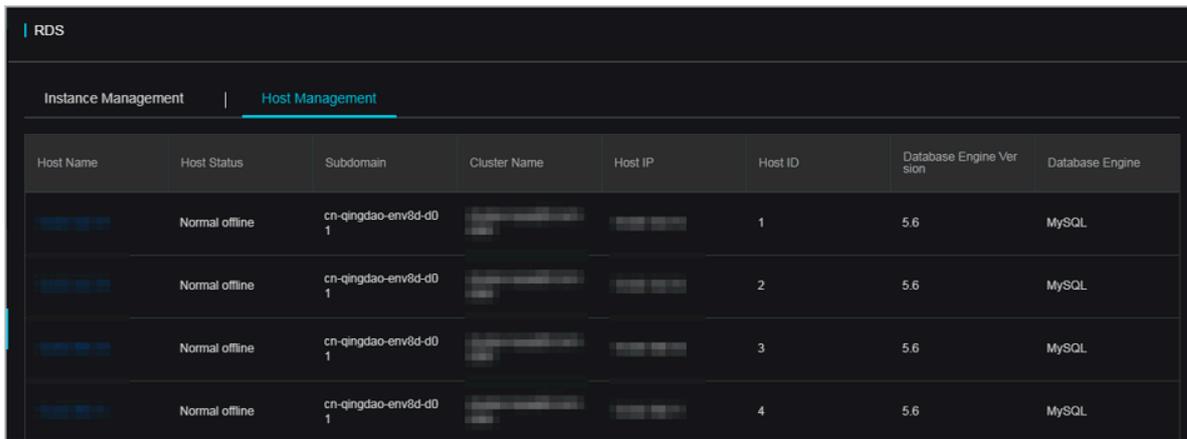
2.7.5 Manage hosts

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**.

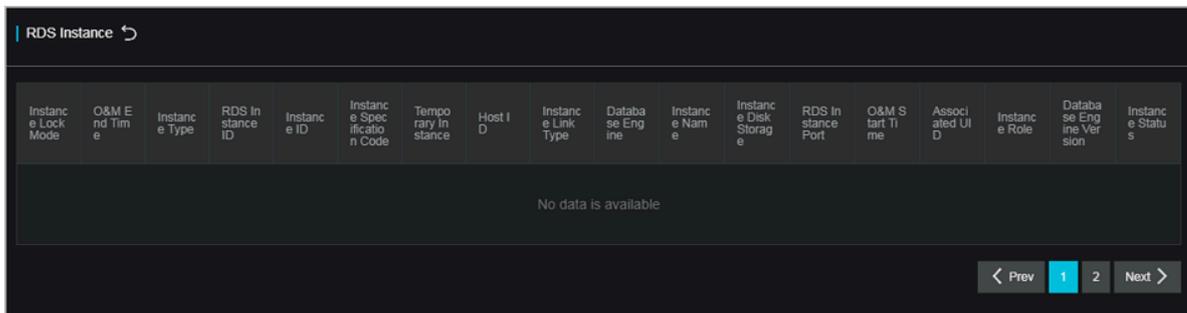
3. On the Host Management tab of RDS, you can view all host information.



The screenshot shows the RDS Host Management interface. It features a table with the following columns: Host Name, Host Status, Subdomain, Cluster Name, Host IP, Host ID, Database Engine Version, and Database Engine. There are four rows of data, all showing 'Normal offline' status and 'MySQL' as the database engine.

Host Name	Host Status	Subdomain	Cluster Name	Host IP	Host ID	Database Engine Version	Database Engine
cn-qingdao-env8d-d01	Normal offline	cn-qingdao-env8d-d01			1	5.6	MySQL
cn-qingdao-env8d-d02	Normal offline	cn-qingdao-env8d-d01			2	5.6	MySQL
cn-qingdao-env8d-d03	Normal offline	cn-qingdao-env8d-d01			3	5.6	MySQL
cn-qingdao-env8d-d04	Normal offline	cn-qingdao-env8d-d01			4	5.6	MySQL

4. Click a hostname to go to the RDS Instance page. You can view all instances on this host.



The screenshot shows the RDS Instance page. The table has the following columns: Instance Lock Mode, O&M End Time, Instance Type, RDS Instance ID, Instance ID, Instance Specification Code, Temporary Instance, Host ID, Instance Link Type, Database Engine, Instance Name, Instance Disk Storage, RDS Instance Port, O&M Start Time, Associated UID, Instance Role, Database Engine Version, and Instance Status. The table is currently empty, displaying 'No data is available'.

Instance Lock Mode	O&M End Time	Instance Type	RDS Instance ID	Instance ID	Instance Specification Code	Temporary Instance	Host ID	Instance Link Type	Database Engine	Instance Name	Instance Disk Storage	RDS Instance Port	O&M Start Time	Associated UID	Instance Role	Database Engine Version	Instance Status
No data is available																	

2.7.6 Security maintenance

2.7.6.1 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.

2.7.6.2 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.

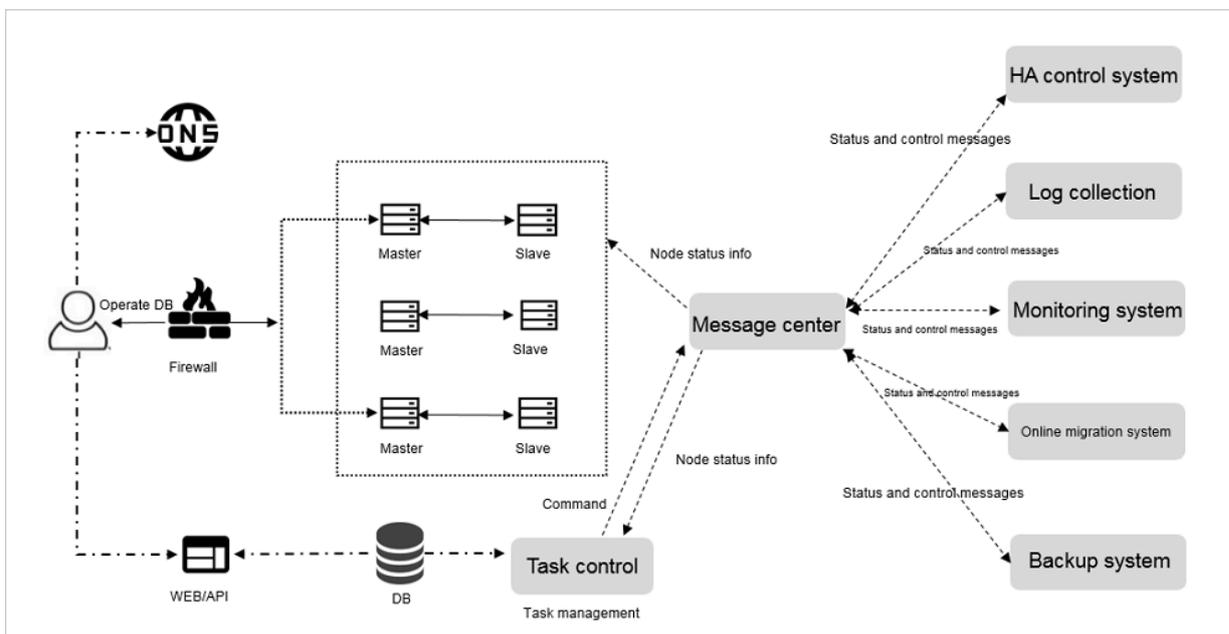
2.8 KVStore for Redis

2.8.1 O&M tool

The Apsara Stack Operation console provides the following operations and maintenance (O&M) features for KVStore for Redis:

- **Instance management:** allows you to view instance details, instance logs, and user information.
- **Host management:** allows you to view and manage hosts.

2.8.2 Architecture diagram



2.8.3 Log on to Apsara Stack Operations

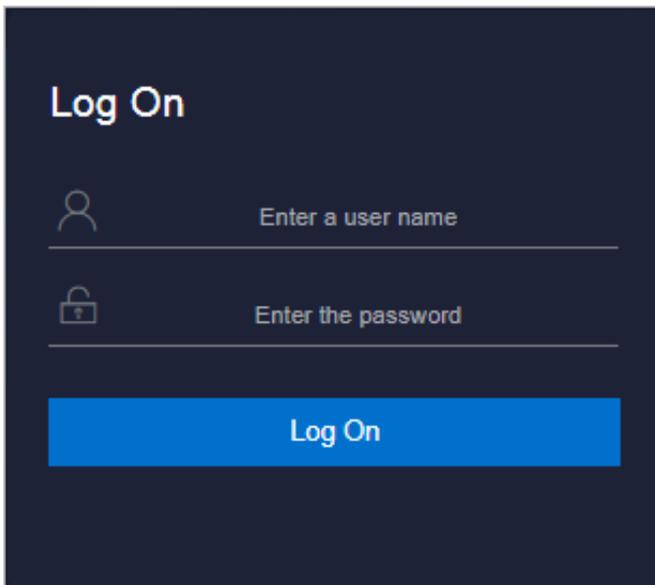
Prerequisites

- **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.

Figure 2-14: 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.

2.8.4 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:

- 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.

2.8.5 Host management

Host management allows you to view and manage hosts.

Procedure

1. *Log on to Apsara Stack Operations.*

- 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.

The screenshot shows the RDS Host Management page. At the top, there are two tabs: 'Instance Management' and 'Host Management', with 'Host Management' selected. Below the tabs is a table with the following columns: Host Name, Host Status, Subdomain, Cluster Name, Host IP, Host ID, Database Engine Version, and Database Engine. The table contains several rows of data, each representing a host. At the bottom of the page, there is a copyright notice: '© 2009-2018 Alibaba Cloud Computing Limited. All rights reserved.'

- Click a host name to go to the RDS Instance page. You can view all instances on this host.

The screenshot shows the RDS Instance page. At the top, there is a breadcrumb 'RDS Instance' with a back arrow. Below it is a table with the following columns: Instance Lock Mode, O&M End Time, Instance Type, RDS Instance ID, Instance ID, Instance Specific Code, Tempo Instance, Host ID, Instance Link Type, Datab... Engine, Instance Name, Instance Disk Storage, RDS Instance Port, O&M Start Time, Instance Role, Datab... Engine Version, and Instance Status. The table contains several rows of data, each representing an instance. At the bottom of the page, there is a copyright notice: '© 2009-2018 Alibaba Cloud Computing Limited. All rights reserved.'

2.8.6 Security maintenance

2.8.6.1 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.

2.8.6.2 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.

2.9 ApsaraDB for MongoDB

2.9.1 Service architecture

2.9.1.1 System architecture

2.9.1.1.1 Backup system

Automatic backup

ApsaraDB for MongoDB supports both physical backup and logical backup.

You can flexibly configure the backup start time based on the service off-peak hours. All backup files are retained for seven days.

Temporary backup

You can initiate a temporary backup as required. The backup files are retained for seven days.

Log management

ApsaraDB for MongoDB generates operation logs and allows you to download them. You can use the operation logs for local incremental backup.

Data backtracking

ApsaraDB for MongoDB can use backup files and logs to generate a temporary instance for any time point within the past seven days. After verifying that the data in the temporary instance is correct, you can use the temporary instance to restore data to the specified time point.

Creating a temporary instance does not affect the running of the current instance.

Only one temporary instance can be created for each ApsaraDB for MongoDB instance at a time. A temporary instance is valid for 48 hours. You can create a maximum of 10 temporary instances for an ApsaraDB for MongoDB instance each day.

2.9.1.1.2 Data migration system

Database replication between instances

ApsaraDB for MongoDB allows you to easily migrate databases from one instance to another.

Data migration to or from ApsaraDB for MongoDB

ApsaraDB for MongoDB provides a professional tool and a migration wizard to help you migrate data to or from ApsaraDB for MongoDB.

Backup file download

ApsaraDB for MongoDB retains backup files for seven days. During this period, you can log on to the ApsaraDB for MongoDB console to download the backup files.

2.9.1.1.3 Monitoring system

Performance monitoring

ApsaraDB for MongoDB provides nearly 20 metrics for system performance monitoring, such as disk capacity, IOPS, connections, CPU utilization, network traffic, TPS, QPS, and cache hit rate. You can obtain the running status information of instances over the past year.

SQL auditing

The system records the SQL statements and related information sent to ApsaraDB for MongoDB instances, such as the connection IP address, database name, access account, execution time, and the number of records returned. You can use SQL auditing to locate problems and check instance security.

Threshold alarms

ApsaraDB for MongoDB provides alarm SMS notifications in the event of exceptions in instance status or performance.

These exceptions can include instance locking, insufficient disk capacity, abnormal IOPS, abnormal number of connections, and over high CPU utilization. You can configure alarm thresholds and up to 50 alarm contacts (of which only five are effective at a time). When a threshold is exceeded in an instance, an SMS notification is sent to the alarm contacts.

Web operation logs

The system logs all modification operations in the ApsaraDB for MongoDB console for administrators to check. These logs are retained for a maximum of 30 days.

2.9.1.1.4 Control system

If a host or an instance crashes, the ApsaraDB for MongoDB high-availability (HA) component fails services over within 30 seconds after the exception is detected. This guarantees that applications run properly and ApsaraDB for MongoDB is highly available.

2.9.1.1.5 Task scheduling system

You can use the ApsaraDB for MongoDB console or APIs to create or delete instances or switch instances between the intranet and Internet. All instance operations are scheduled, traced, and displayed as tasks.

2.9.2 ApsaraDB for MongoDB O&M overview

Apsara Stack Operations Console provides the following O&M features for ApsaraDB for MongoDB:

- **Instance management:** allows you to view instance details, instance logs, and user information.
- **Host management:** allows you to view and manage hosts.

2.9.3 Log on to the Apsara Stack Operations console

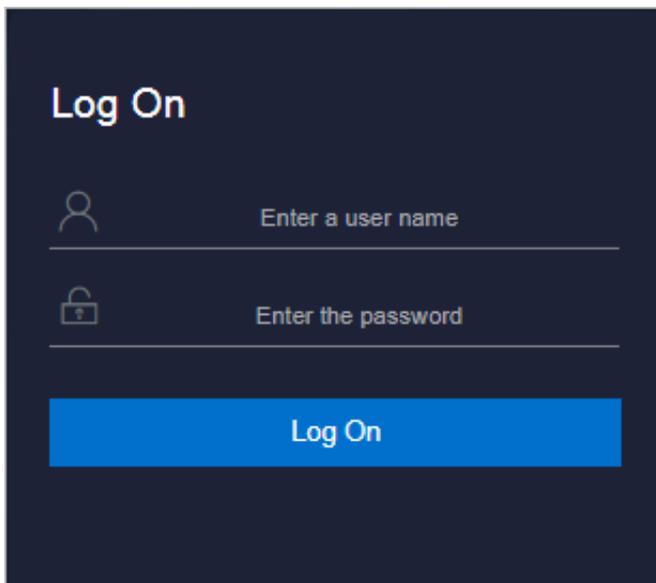
Prerequisites

- 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.

Figure 2-15: 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.

2.9.4 Manage instances

You can view instance details, logs, and user information.

Procedure

1. *Log on to the Apsara Stack Operations console.*

2. In the left-side navigation pane, choose **Products > RDS** to go to the RDS page. Click the **Instance Management** tab. On the Instance Management tab, you can perform the following operations:

- **Instances**

View the instances that belong to the account on the Instance Management tab, as shown in *Figure 2-16: Instances*.

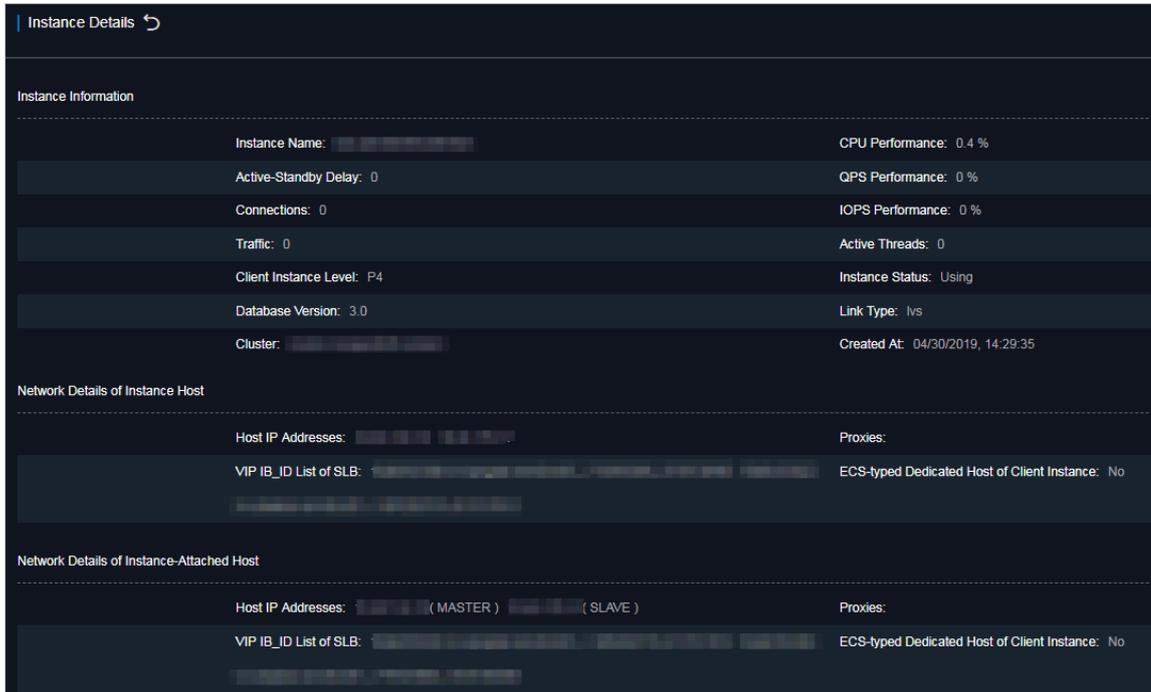
Figure 2-16: Instances

Instance Name	Availa...	CPU Perfor...	QPS Perfor...	IOPS Perfor...	Conne...	Disk Usage	Instance Status	Database Type	Actions
[REDACTED]	Yes	1.6 %			0	0.5 %	Using	gpdb	User Information
[REDACTED]	Yes	1.5 %			4.5		Using	ppassql	User Information
[REDACTED]	Yes	1.4 %			4	0.1 %	Using	ppassql	User Information
[REDACTED]	Yes	1.4 %			0	0.3 %	Using	gpdb	User Information
[REDACTED]	Yes	1.2 %			0	0.2 %	Using	gpdb	User Information
[REDACTED]	Yes	1.11 %	5.62 %		0	1.3 %	Using	mysql	User Information

- **Instance details**

Click the ID of an instance to view details, as shown in [Figure 2-17: Instance details](#).

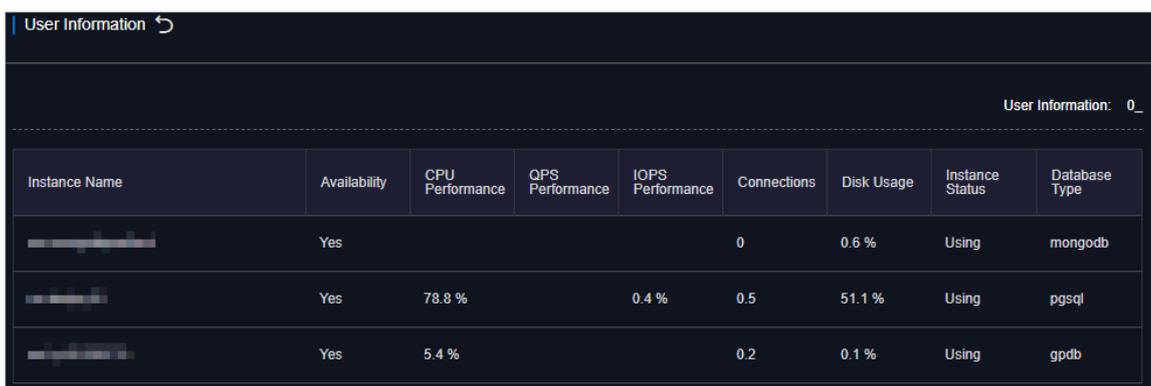
Figure 2-17: Instance details



- **User information**

Click **User Information** in the **Actions** column, as shown in [Figure 2-18: User information](#).

Figure 2-18: User information



2.9.5 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.

Host Name	Host Status	Subdomain	Cluster Name	Host IP	Host ID	Database Engine Version	Database Engine
[Redacted]	normal operation	[Redacted]	[Redacted]	[Redacted]	1	5.6	MySQL
[Redacted]	normal operation	[Redacted]	[Redacted]	[Redacted]	2	5.6	MySQL
[Redacted]	normal operation	[Redacted]	[Redacted]	[Redacted]	3	5.6	MySQL
[Redacted]	normal operation	[Redacted]	[Redacted]	[Redacted]	4	5.6	MySQL
[Redacted]	normal operation	[Redacted]	[Redacted]	[Redacted]	5	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.

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
0	06:00	Primary Instance	417	[Redacted]	dds.mongo.mid	No	7	ivs	MongoDB
0	06:00	Primary Instance	420	[Redacted]	dds.mongo.mid	No	7	ivs	MongoDB
0	06:00	Primary Instance	1546	[Redacted]	dds.mongo.mid	No	7	ivs	MongoDB

2.9.6 Security maintenance

2.9.6.1 Network security maintenance

Network security maintenance is aimed at ensuring device security and network security.

Device security

Check network devices, and enable security management protocols and configurations of devices.

Check for up-to-date versions of network device software and update the software to more secure versions in a timely manner.

For more information about the security maintenance method, see the product document of each device.

Network security

Select the intrusion detection system (IDS) or intrusion prevention system (IPS) based on the network status to check Internet and intranet traffic and defend the network against abnormal behaviors and attacks.

2.9.6.2 Account password maintenance

Account passwords include the ApsaraDB for MongoDB system and device passwords.

To ensure account security, change the system and device passwords periodically, and use passwords that meet the complexity requirements.

2.10 AnalyticDB for PostgreSQL

2.10.1 Overview

Purpose

This guide summarizes possible problems that you may encounter during O&M operations and provides solutions for you.

If you encounter system problems not covered in this guide, you can submit a ticket to Alibaba Cloud for technical support.

Requirements

You must possess IT skills, including computer network knowledge, computer operation knowledge, and capabilities of problem analysis and troubleshooting.

In addition, 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 operations on the system, you must be aware of the results caused by operations to avoid the impact on system operation.

You must record all problems encountered during operations for problem analysis and troubleshooting.

- **Personal and data security**

- **You must take safety measures according to device manuals when operating electrical equipment.**
- **You must use secure devices to access the business network.**
- **Unauthorized data replication and dissemination are prohibited.**

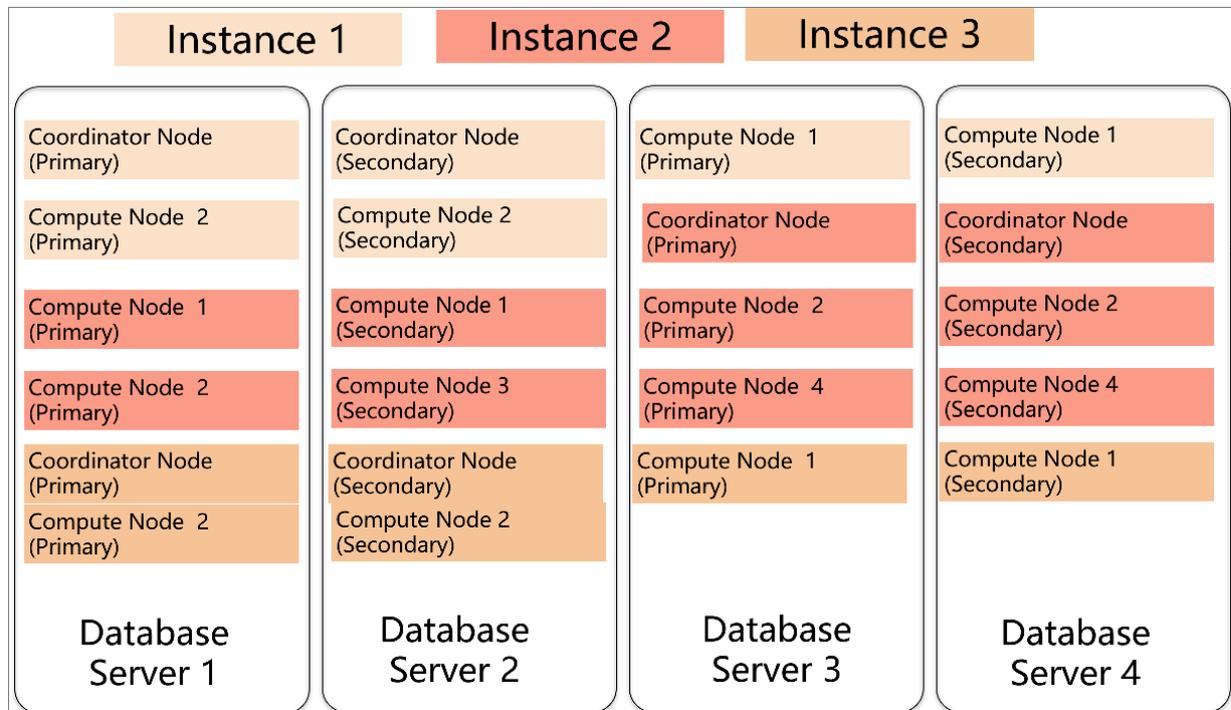
Support

You can contact Alibaba Cloud technical support for help.

2.10.2 System architecture

The following figure shows the AnalyticDB for PostgreSQL architecture.

Figure 2-19: System architecture



AnalyticDB for PostgreSQL consists of two major components: the master and segment nodes.

The master is used to access applications. The master accepts connection requests and SQL query requests from clients and dispatches computing tasks to segments. The system deploys a secondary node of the master on an independent physical server and replicates data from the primary node to the secondary node for failover. The secondary node cannot connect to segments or accept external links.

Segments are independent data nodes in AnalyticDB for PostgreSQL. Each segment stores a part of the data, and all the segments execute computing tasks at the same time to implement parallel processing. Each segment consists of a primary node and a secondary node for failover.

2.10.3 Routine maintenance

2.10.3.1 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 groups.

You can use the following methods to locate the table skew. The procedure is as follows.

1. For a single table or database, you can view the space occupied in each segment to check whether the data skew occurs.

a. Execute the following statement to check whether the data in a database is 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 segment after the statement is executed. If the space occupied in one or more segments is significantly greater than that of other segments, it indicates the data in this database is skewed.

b. Execute the following statement to check whether the data in a table is skewed.

```
select pg_size_pretty(pg_relation_size('tblname')) from gp_dist_ra
ndom('gp_id');
```

You can view the space occupied by the tblname table in each segment after the statement is executed. If the space occupied in one or more segments is significantly greater than that of other segments, 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 check whether data skew occurs.
 - a. Execute the following statement to check whether the storage space is skewed. The principle is similar to that of the 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 takes for the check to complete.

- b. You can 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](#).

2.10.3.2 Execute the VACUUM and ANALYZE statements

You can execute `VACUUM` and `ANALYZE` statements for frequently updated tables and databases on a regular basis. You can also execute `VACUUM` and `ANALYZE` statements after you have performed a large number of update or write operations, preventing these operations from consuming excessive system resources and storage.

2.10.4 Security maintenance

2.10.4.1 Network security maintenance

Network security maintenance helps you ensure 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

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.

2.10.4.2 Account password maintenance

The 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 change the system and device passwords periodically.

2.11 Apsara Stack Security

2.11.1 Log on to the Apsara Infrastructure Management Framework console

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

Prerequisites

Before logging on to the console, obtain the URL of Apsara Stack Operations Console, and the username and password to log on to the console from your system administrator.

Procedure

1. In the address bar of a browser, enter `https://ASOP URL`, and press Enter.
2. On the logon page, enter the username and password, and then click Log On.
3. In the left-side navigation pane, select `Products`.
4. In the product list, click `Apsara Infrastructure Management Framework` to go to the Apsara Infrastructure Management Framework console.

2.11.2 Routine operations and maintenance of Server Guard

2.11.2.1 Check the service status

2.11.2.1.1 Check the client status

Check the following status information about the Server Guard client to verify that the client is running properly:

Client logs

Client logs are stored in the data directory under the directory of the Server Guard process file, for example, `/usr/local/aegis/aegis_client/aegis_xx_xx/data`.

Client logs are saved by day, for example, `data.1` to `data.7`

Client's online status

Run the following command to check the client's online status:

```
ps -aux | grep AliYunDun
```

Network connectivity

Run the following command to check whether the client has set up a TCP connection with the server:

```
netstat -tunpe |grep AliYunDun
```

Client UUID

Open the client log file `data.x` and check the character string following `Currentuid` Ret. This character string is the UUID of the current client.

Client processes

The Server Guard client has three resident processes: `AliYunDun`, `AliYunDunUpdate`, and `AliHids`.

When the client runs properly, all of the three processes run normally.



Note:

On a Windows OS client, the `AliYunDun` and `AliYunDunUpdate` processes exist in the form of services. The service names are `Server Guard Detect Service` and `Server Guard Update Service`, respectively.

2.11.2.1.2 Check the status of Aegiserver

Context

To check the running status of Aegiserver, follow the following steps:

Procedure

- 1. Run the `ssh server IP address` command to log on to the server of Aegiserver.**
- 2. Run the following command to find the Aegiserver image ID:**

```
docker ps -a |grep aegiserver
```

The following message is displayed:

```
b9e59994df41  
reg.docker.alibaba-inc.com/aqs/aegiserverlite@sha256:f9d292f54c  
58646b672a8533a0d78fba534d26d376a194034e8840c70d9aa0b3 "/bin/bash /
```

```
startApp." 2 hours ago Up 2 hours 80/tcp, 7001/tcp, 8005/tcp, 8009/tcp
yundun-aegis.Aegiserverlite_..aegiserverlite. 1484712802
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to check whether the Java process is normal:

```
ps aux |grep aegiserver
```

The following message is displayed:

```
root 153 0.6 25.8 2983812 1084588 ? Sl 12:13 1:01 /opt/taobao/java
/bin/java -Djava.util.logging.config.file=/home/admin/aegiserver
lite/.default/conf/logging.properties -Djava.util.logging.manager
=org.apache.juli.ClassLoaderLogManager -server -Xms2g -Xmx2g -XX:
PermSize=96m -XX:MaxPermSize=384m -Xmn1g -XX:+UseConcMarkSweepGC -XX
:+UseCMSCompactAtFullCollection -XX:CMSMaxAbortablePrecleanTime=5000
-XX:+CMSClassUnloadingEnabled -XX:+UseCMSInitiatingOccupancyOnly -XX
:CMSInitiatingOccupancyFraction=80 -XX:+HeapDumpOnOutOfMemoryError -
XX:HeapDumpPath=/home/admin/logs/java.hprof -verbose:gc -Xloggc:/home
/admin/logs/gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -Djava
.awt.headless=true -Dsun.net.client.defaultConnectTimeout=10000 -Dsun
.net.client.defaultReadTimeout=30000 -XX:+DisableExplicitGC -Dfile.
encoding=UTF-8 -Ddruid.filters=mergeStat -Ddruid.useGlobalDataSourceSt
at=true -Dproject.name=aegiserverlite -Dcatalina.vendor=alibaba -Djava
.security.egd=file:/dev/./urandom -Dlog4j.defaultInitOverride=true -
Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_EQUALS_IN_VALUE=true -
Dorg.apache.tomcat.util.http.ServerCookie.ALLOW_HTTP_SEPARATORS_IN_V0=
true -Djava.endorsed.dirs=/opt/taobao/tomcat/endorsed -classpath /opt/
taobao/tomcat/bin/bootstrap.jar:/opt/taobao/tomcat/bin/tomcat-juli.jar
-Dcatalina.logs=/home/admin/aegiserverlite/.default/logs -Dcatalina.
base=/home/admin/aegiserverlite/.default -Dcatalina.home=/opt/taobao/
tomcat -Djava.io.tmpdir=/home/admin/aegiserverlite/.default/temp org.
apache.catalina.startup.Bootstrap -Djboss.server.home.dir=/home/admin
/aegiserverlite/.default -Djboss.server.home.url=file:/home/admin/
aegiserverlite/.default start
```

5. Run the following command to perform the health check:

```
curl 127.0.0.1:7001/checkpreload.htm
```

If the response is "success", the service is normal.

6. View related logs.

- **Protocol logs:** View logs about upstream and downstream protocol messages between the server and client in `/home/admin/aegiserver/logs/AEGIS_MESS`
`AGE.log`.
- **Operation logs:** View abnormal stack information during operation in `/home/`
`admin/aegiserver/logs/aegis-default.log`.
- **Offline logs:** View the logs about client disconnection caused by time-out in `/`
`home/admin/aegiserver/logs/AEGIS_OFFLINE_MESSAGE.log`.

2.11.2.1.3 Check the Server Guard Update Service status

Context

To check the status of Server Guard Update Service, follow the following steps:

Procedure

1. Run the `ssh host IP address` command to log on to the server of Aegiserver.
2. Run the following command to find the Aegiserver image ID:

```
docker ps -a |grep aegiserver
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to check whether the Java process is normal:

```
ps aux |grep aegisupdate
```

5. Run the following command to perform the health check:

```
curl 127.0.0.1:7001/checkpreload.htm
```

If the response is "success", the service is normal.

2.11.2.1.4 Check the Defender module status

Context

To check the status of the Defender module of Server Guard, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the Defender module of Server Guard.
2. Run the following command to find the image ID of the Defender module of Server Guard:

```
docker ps -a |grep defender
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to check whether the Java process is normal:

```
ps aux |grep defender
```

5. Run the following command to perform health check:

```
curl 127.0.0.1:7001/checkpreload.htm
```

If the response is "success", the service is normal.

2.11.2.2 Restart Server Guard

Context

To restart Server Guard when a fault occurs, follow these steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts Server Guard.

2. Run the following command to find the image ID of Server Guard:

```
docker ps -a |grep application name
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Restart related services.

- **Restart the Server Guard client service.**
 - **For a server running a Windows OS, go to the service manager, locate *Server Guard Detect Service*, and restart this service.**
 - **For a server running a Linux OS, use either of the following methods to restart the Server Guard client service:**
 - **Run the `service aegis restart` command to restart the service.**
 - **Run the `killall AliYunDun` command as the root user to stop the current process, and then restart the `/usr/local/aegis/aegis_client/aegis_xx_xx/AliYunDun` process.**
- **Restart the Aegiserver service.**
 - a. **Run the following command to view the Java process ID:**

```
ps aux |grep aegiserver
```
 - b. **Run the following command to stop the current process:**

```
kill -9 process
```
 - c. **Run the following command to restart the process:**

```
sudo -u admin /home/admin/aegiserver/bin/jbossctl restart
```
 - d. **Run the following command to check whether the process has been successfully restarted:**

```
curl 127.0.0.1:7001/checkpreload.htm
```
- **Restart Server Guard Update Service:**
 - a. **Run the following command to view the Java process ID:**

```
ps aux |grep aegisupdate
```
 - b. **Run the following command to stop the current process:**

```
kill -9 process
```
 - c. **Run the following command to restart the process:**

```
sudo -u admin /home/admin/aegisupdate/bin/jbossctl restart
```
 - d. **Run the following command to check whether the process has been successfully restarted:**

```
curl 127.0.0.1:7001/checkpreload.htm
```

- **Restart the Defender service of Server Guard.**
 - a. **Run the following command to view the Java process ID:**

```
ps aux |grep secure-service
```

- b. **Run the following command to stop the current process:**

```
kill -9 process
```

- c. **Run the following command to restart the process:**

```
sudo -u admin /home/admin/secure-service/bin/jbossctl restart
```

- d. **Run the following command to check whether the process has been successfully restarted:**

```
curl 127.0.0.1:7001/checkpreload.htm
```

2.11.3 Routine operations and maintenance of Network Traffic Monitoring System

2.11.3.1 Check the service status

2.11.3.1.1 Basic inspection

During the basic inspection of Network Traffic Monitoring System, check whether the service has reached the final status.

Procedure

1. *Log on to the Apsara Infrastructure Management Framework console.*
2. **Choose Operations > Project Operations. On the page that appears, enter yundun-advance, and click Details to go to the Cluster Operations page.**
3. **Select BasicCluster.**
4. **Check whether yundun-beaver-advance has reached the final status in Service Instances List.**

2.11.3.1.2 Advanced inspection

During the advanced inspection feature of Network Traffic Monitoring System, check the status and features of the service.

Procedure

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

2. Log on to two physical machines of Network Traffic Monitoring System, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the **Cluster Operations** page.
 - c) Select **BasicCluster**.
 - d) Select `yundun-beaver-advance` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select `BeaverAdvance#` from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **Server Information**, and use **TerminalService** to log on to two physical machines of Network Traffic Monitoring System, respectively.
3. Check the log status of Network Traffic Monitoring System.

Run `sudo cat /var/log/messages`. If any record is returned, the logs are normal.
4. Check the status of the mirrored traffic.

Run `sudo cat /proc/ixgbe_debug_info`. If the speed is not 0 in the second-to-last row of the output, the mirrored traffic is normal.
5. Check the configuration of the protected IP CIDR block.

Run `tail -f /dev/shm/banff-2018-xx.log`. In the command, `xx` indicates the month. For example, the log file for May in 2018 is named `banff-2018-05.log`. The IP CIDR block in the output should be the classic network SLB/EIP CIDR block (for CSW non-standard access, configure the VPC CIDR block).
6. Check the network connectivity between Network Traffic Monitoring System and the VM.

Run `ping VMIP` to check the network connectivity. In the command, `VMIP` is a real IP address that falls in the CIDR block of the previous step.
7. Check the `tcp_decode` process status.

Run `ps -ef | grep tcp_decode`. If any record is returned, the `tcp_decode` process is normal.

8. Check the configuration of the traffic scrubbing server.

Run `cat /home/admin/beaver-dj-schedule/conf/dj.conf` and check whether the IP address specified in the unmarked configuration item `aliguard_smart` is the DNS VIP of the domain name `aliguard.${global:internet-domain}`.

9. Check the following typical logs:

- DDoS alert logs

Run the `grep -A 10 -B 10 LIDS /var/log/messages` command to view the DDoS alert logs.

- TCP blocking command logs

Run the `grep add_to_blacklist.htm /var/log/messages` command to view the TCP blocking command logs.

- Outbound attack logs

Run the `grep zombie_new /var/log/messages` command to view the outbound attack logs.

2.11.3.2 Common operations and maintenance

2.11.3.2.1 Restart the Network Traffic Monitoring System process

Context

To restart the Network Traffic Monitoring System process, follow the following steps:

Procedure

1. Log on to the physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to restart the Network Traffic Monitoring System process:

```
rm -rf /dev/shm/drv_setup_path
```

2.11.3.2.2 Uninstall Network Traffic Monitoring System

Context

To uninstall Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to uninstall Network Traffic Monitoring System:

```
bash /opt/beaver/bin/uninstall.sh
```

2.11.3.2.3 Disable TCP blocking

Context

To disable TCP blocking for Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Open the `/beaver_client.sh` file on each server of Network Traffic Monitoring System, and add a number sign (#) to the start of the `./tcp_reset` line to comment out the line.
4. Run the following command on each server of Network Traffic Monitoring System to disable TCP blocking:

```
killall tcp_reset
```

2.11.3.2.4 Enable TCPDump

Context

To enable TCPDump for Network Traffic Monitoring System, follow the following steps:

Procedure

1. Log on to a physical machine of Network Traffic Monitoring System.
2. Switch to the root account.
3. Run the following command to enable TCPDump:

```
echo 1 > /proc/ixgbe_debug_dispatch
```



Note:

When TCPDump is enabled, the performance of Network Traffic Monitoring System may be affected. We recommend that you run the following command to disable TCPDump after packet capture is complete.

```
echo 0 > /proc/ixgbe_debug_dispatch
```

2.11.4 Routine operations and maintenance of Anti-DDoS Service

2.11.4.1 Check the service status

2.11.4.1.1 Basic inspection

The basic inspection of Anti-DDoS Service checks whether the service has reached the final status.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console, and choose **Operations > Project Operations**. Enter `yundun-advance`, and click **Details** to go to the Cluster Operations page.
2. Select **AliguardCluster**.
3. Check whether `yundun-aliguard` has reached the final status in **Service Instances List**.

2.11.4.1.2 Advanced inspection

The advanced inspection of Anti-DDoS Service checks the status and features of the service.

Procedure

To check the running status of Anti-DDoS Service, follow the following steps:

1. Log on to two physical machines of Anti-DDoS Service, respectively.
 - a) Log on to the Apsara Infrastructure Management Framework console, and choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the **Cluster Operations** page.
 - c) Select **AliguardCluster**.
 - d) Select `yundun-aliguard` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select `AliguardConsole#` from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **Server Information**, and use **TerminalService** to log on to two physical machines of Anti-DDoS Service, respectively.
2. Check the deployment status of Anti-DDoS Service.

Run `/home/admin/aliguard/target/AliguardDefender/bin/aliguard_defender_check`, and check the output result.



Note:

If a server of Anti-DDoS Service has just restarted, wait for three to five minutes before running the script to check the deployment status.

- If the message `aliguard status check OK!` appears, Anti-DDoS Service has been correctly deployed and the service status is normal, as shown in [Figure 2-20: Check the status of Anti-DDoS Service](#).

Figure 2-20: Check the status of Anti-DDoS Service

```
1 [root@192.168.1.100.cloud.tencent.com /home/admin]
2 #aliguard_defender_check
3 myfwd
4 aliguard_log
5 netframe
6 route_monitor
7 neigh_monitor
8 aliguard_monitor
9 bgpd
10 rsyslogd
11 aliguard status check OK!
```

- If the error message shown in [Figure 2-21: ReInjection route error message](#) appears, the reInjection route is faulty.

Figure 2-21: ReInjection route error message

```
1 Error: route status error, we need two default routes to reinject the net flow!
2 Error: route error, can't get to the target ip.
```

Troubleshooting: The reInjection route is a default route generated by Anti-DDoS Service and is redirected to the interface through which the ISW is bound to the VPN in the next hop. If any problem occurs, check whether this route has been generated by Anti-DDoS Service. If this route has been

generated, check whether the ISW has forwarded this route to downstream devices.

- If the error message shown in *Figure 2-22: BGP routing error message* appears, the BGP protocol (for traffic routing) is faulty.

Figure 2-22: BGP routing error message

1 Error: bgp status error!

Troubleshooting: If BGP routing is faulty, troubleshoot the problem as follows:

- Use the ISW to check whether the BGP neighbor is in the normal status.
 - 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.
 - Check whether the route policy in the BGP configuration of the ISW is correctly configured.
- If the problem is caused by none of the above reasons, the core process is faulty. Contact Alibaba Cloud technical support.
3. Check the status of the NICs or optical modules of Anti-DDoS Service.



Note:

Anti-DDoS Service has special requirements on optical modules. Only optical modules equipped with Intel X520 or Intel 82599 NICs can be used.

Run `lspci | grep Eth`. If the command output contains four Intel 82599 NICs, the NICs are standard.

```
[root@cloud.am54 /root]
#lspci -v | grep Eth
02:00.0 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
02:00.1 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
04:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
    Subsystem: Intel Corporation Ethernet Server Adapter X520-2
04:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
    Subsystem: Intel Corporation Ethernet Server Adapter X520-2
81:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
    Subsystem: Intel Corporation Ethernet Server Adapter X520-2
81:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit SFI/SFP+ Network Connection (rev 01)
    Subsystem: Intel Corporation Ethernet Server Adapter X520-2
```

2.11.4.2 Common operations and maintenance

2.11.4.2.1 Restart Anti-DDoS Service

Context

To restart Anti-DDoS Service when an error occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts Anti-DDoS Service.
2. Run the following command to stop Anti-DDoS Service:

```
/home/admin/aliguard/target/AliguardDefender/bin/aliguard stop
```



Note:

If the `ERROR: Module net_msg is in use` message is displayed, run the command again later. If Anti-DDoS Service cannot be stopped after several attempts, restart the server of Anti-DDoS Service.

3. Run the following command to restart Anti-DDoS Service:

```
/home/admin/aliguard/target/AliguardDefender/bin/aliguard start
```

4. Run the service status check command five minutes after Anti-DDoS Service is restarted.

2.11.4.2.2 Troubleshoot common faults

Context

When an error occurs in Anti-DDoS Service, follow the following troubleshooting steps:

Procedure

1. Restart Anti-DDoS Service.

- If Anti-DDoS Service is in the normal status after being restarted but an error message is returned during the health check performed later, non-standard NICs or optical modules are used. To check whether standard NICs or optical modules are used, see [Check the status of the NICs or optical modules of Anti-DDoS Service](#). If non-standard NICs or optical modules are used, change the NICs or optical modules.
- If Anti-DDoS Service is in an unusual status after being restarted, go to the next step.

2. View the `aliguard_dynamic_config` file.

Carefully check whether each configuration item in the file is exactly the same as that in the plan.



Note:

Ensure that the AS number specified in `aliguard local` is 65515 and that the BGP password is correct.

3. Check the wiring and switch configuration.



Note:

If any incorrect configuration is found, the current fault is caused by incorrect wiring or switch IP address configuration, rather than incorrect deployment of Anti-DDoS Service. In this case, contact the network engineer.

Assume that the Anti-DDoS Service configurations to be checked are listed in the following figure, among which the server IP address is 10.1.4.12. To check

whether the four ports of Anti-DDoS Service can ping the ports of the switch, follow the following steps:

Figure 2-23: Anti-DDoS Service configuration example

aliguard_host_ip	port	aliguard_port_ip	csr_port_ip
10.1.4.12	T0	10.1.0.34	10.1.0.33
10.1.4.12	T1	10.1.0.38	10.1.0.37
10.1.4.12	T2	10.1.0.50	10.1.0.49
10.1.4.12	T3	10.1.0.54	10.1.0.53
10.1.4.28	T0	10.1.0.42	10.1.0.41
10.1.4.28	T1	10.1.0.46	10.1.0.45
10.1.4.28	T2	10.1.0.58	10.1.0.57
10.1.4.28	T3	10.1.0.62	10.1.0.61

a. Run the following commands to check the NIC PCI IDs of Anti-DDoS Service:

```
cd /sys/bus/pci/drivers/igb_uio
```

```
ls
```

Record the PCI IDs of the four NICs, for example, 0000:01:00.0, 0000:01:00.1, 0000:82:00.0, and 0000:82:00.1.

b. Run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard stop` command to stop Anti-DDoS Service.

c. In the `/sys/bus/pci/drivers/igb_uio` directory, unbind the four NICs recorded in the first step from the `igb_uio` driver, as shown in [Figure 2-24: Unbind NICs](#).

Figure 2-24: Unbind NICs

```
1 echo "0000:01:00.0" >> unbind
2 echo "0000:01:00.1" >> unbind
3 echo "0000:82:00.0" >> unbind
4 echo "0000:82:00.1" >> unbind
```

- d. In the `/sys/bus/pci/drivers/ixgbe` directory, bind the four NICs to the `ixgbe` driver for Linux, as shown in [Figure 2-25: Bind NICs](#).

Figure 2-25: Bind NICs

```
1 echo "0000:01:00.0" >> bind
2 echo "0000:01:00.1" >> bind
3 echo "0000:82:00.0" >> bind
4 echo "0000:82:00.1" >> bind
```

- e. Set Anti-DDoS Service IP addresses for the NICs.

The local server IP address is 10.1.4.12, and the NIC IP addresses are set to 10.1.0.34, 10.1.0.38, 10.1.0.50, and 10.1.0.54, as shown in [Figure 2-23: Anti-DDoS Service configuration example](#).

- A. Run the `ifconfig-a` command to display all NICs, and run the `ethtool -i` command to view the PCI ID of each NIC. Find the four NICs of which the

IDs are the same as those recorded in the first step, for example, eth0, eth1, eth2, and eth3.

B. Run the following commands to move these NICs to the top of the queue:

```
ifconfig eth0 up
```

```
ifconfig eth1 up
```

```
ifconfig eth2 up
```

```
ifconfig eth3 up
```

C. Set Anti-DDoS Service IP addresses for the NICs. Run the following commands to set Anti-DDoS Service IP addresses for the NICs based on their PCI IDs in an ascending order:

```
ifconfig eth0 10.1.0.34 netmask 255.255.255.252
```

```
ifconfig eth1 10.1.0.38 netmask 255.255.255.252
```

```
ifconfig eth2 10.1.0.50 netmask 255.255.255.252
```

```
ifconfig eth3 10.1.0.54 netmask 255.255.255.252
```

f. Try to ping the peer IP addresses configured. If the peer IP addresses cannot be pinged, the switch configuration or wiring is incorrect.

```
ping 10.1.0.33
```

```
ping 10.1.0.37
```

```
ping 10.1.0.49
```

```
ping 10.1.0.53
```

g. If these four IP addresses can all be pinged, you can directly start Anti-DDoS Service without unbinding the NICs.

Run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard start` **command to start Anti-DDoS Service.**

After Anti-DDoS Service has been started for a while, run the `/home/admin/aliguard/target/AliguardDefender/bin/aliguard_rule -v 0.0.0.0 -d drop_icmp` **command to disable the drop_icmp policy.**

h. Ping the peer IP addresses again.

```
ping 10.1.0.33
```

```
ping 10.1.0.37
```

```
ping 10.1.0.49
```

```
ping 10.1.0.53
```

If the peer IP addresses cannot be pinged, non-standard NICs or optical modules are used or the configuration is incorrect.

- 4. If these four peer IP addresses can be pinged after Anti-DDoS Service is started but an error is reported during a status check of Anti-DDoS Service, contact Alibaba Cloud technical support.**

2.11.5 Routine operations and maintenance of Threat Detection Service

2.11.5.1 Check the service status

2.11.5.1.1 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.**

2.11.5.1.2 Advanced inspection

The advanced inspection of TDS checks the status and features of the service.

Procedure

To check the TDS running status, follow the following steps:

- 1. Log on to the Apsara Infrastructure Management Framework console.*

2. Log on to two TDS physical machines, respectively.

- a) **Choose Operations > Project Operations.**
- b) **Enter yundun-advance, and click Details to go to the Cluster Operations page.**
- c) **Select BasicCluster.**
- d) **Select yundun-sas from Service Instances List, and click Details to go to the Service Instance Dashboard page.**
- e) **Select SasApp# from Service Role List, and click Details to go to the Service Role Dashboard page.**
- f) **View Server Information, and use TerminalService to log on to two TDS physical machines, respectively.**

3. Log on to two TDS Docker containers, respectively.

Run `sudo docker exec -it $(sudo docker ps | grep sas | awk '{print $1}') bash.`

4. Check the Java process status.

Run `ps aux |grep sas`. If any record is returned, the process is normal.

5. Check the health status.

Run `curl 127.0.0.1:3008/check.htm`. If OK is returned, the service is normal.

6. View related logs.

- **View all logs in `/home/admin/sas/logs/sas-default.log`, including metaq message logs, execution logs of scheduled tasks, and error logs. Typically, you can locate TDS faults based on these logs.**
- **View the info logs generated when TDS is running in `/home/admin/sas/logs/common-default.log`.**
- **View the TDS error logs in `/home/admin/sas/logs/common-error.log`.**
- **View the logs about metaq messages received by TDS in `/home/admin/sas/logs/SAS_LOG.log`.**



Note:

Asset verification has been performed on messages in this log file, and the number of messages in this log file is less than that in the `sas-default.log` file.

- **View the logs generated when the alert contact sends an alert notification in `/home/admin/sas/logs/notify.log`.**

2.11.5.2 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
```

2.11.6 Routine operations and maintenance of WAF

2.11.6.1 Check the service status

2.11.6.1.1 Basic inspection

The basic inspection feature of Web Application Firewall (WAF) focuses on whether the service has reached the final status.

Procedure

1. *Log on to the Apsara Infrastructure Management Framework console.*
2. Choose **Operations > Project Operations**.
3. In the **Fuzzy Search** search box, enter `yundun-semawaf`. The search results are displayed.

4. Click Details in the Actions column. The Cluster Operations page is displayed.
5. In the cluster list, click the cluster name that starts with SemaWafCluster.
6. In the Service Instances area on the Cluster Dashboard page, check whether the yundun-semawaf service instance is in final status.

**Note:**

If the Final Status column for an instance is True, the instance has reached final status.

2.11.6.1.2 Advanced inspection

The advanced inspection feature of Web Application Firewall (WAF) focuses on the system status and service status.

Procedure

1. *Log on to the Apsara Infrastructure Management Framework console.*
2. **Log on to two WAF physical machines respectively.**
 - a) **In Apsara Infrastructure Management Framework, choose Operations > Project Operations.**
 - b) **In the Fuzzy Search search box, enter yundun-semawaf. Click Details in the Actions column, and the Cluster Operations page is displayed.**
 - c) **Click the SemaWafCluster cluster.**
 - d) **In Service Instances, select yundun-semawaf, and click Details. The Service Instance Information Dashboard page is displayed.**
 - e) **In Server Role List, select YundunSemawafApp#, and click Details. The Server Role Dashboard page is displayed.**
 - f) **In Machine Information, click Terminal to log on to two WAF physical machines respectively.**

3. Check the system status.

a) Check the system logs.

Run the `dmesg -T |tail -30` command to check for exception logs.

b) Check the system load.

- Run the `free -h` command to check whether the memory usage is normal.
- Run the `df -h` command to check whether the disk usage is normal.
- Run the `uptime` command to check whether the system load average is normal.
- Run the `top` command to check whether the CPU usage is normal.

4. Check the service status.



Note:

The following check is based on the WAF installation directory, which is `/home/safeline` by default.

- a) Run the `cd /home/safeline` command to open the installation directory.
- b) Check the minion service.
 - A. Run the `systemctl status minion` command to check the execution time and status of the minion service.
 - B. Run the `tail -100 logs/minion/minion.log` command to check for exception logs.
- c) Check the mgt-api service.
 - A. Run the `docker logs --tail 50 mgt-api` command to check for exception logs.
 - B. Run the `docker exec -it mgt-api supervisorctl status` command to check whether the service runs normally and whether uptime is normal.
 - C. Run the `tail -50 logs/management/gunicorn.log` command to check for exception logs.
 - D. Run the `tail -50 logs/management/daphne.log` command to check for exception logs.
 - E. Run the `tail -50 logs/management/scheduler.log` command to check for exception logs.
 - F. Run the `tail -50 logs/management/dramatiq.log` command to check for exception logs.
- d) Check the Redis service.

Run the `docker logs --tail 50 mgt-redis` command to check for exception logs.
- e) Check the detector service.
 - A. Run the `docker logs --tail 50 detector-srv` command to check for exception logs.
 - B. Run the `tail -50 logs/detector/snserver.log` command to check for exception logs.
 - C. Run the `curl 127.0.0.1:8001/stat | grep num` command to check whether the service responds normally and whether the real-time request processing data is normal. For example, check the `req_num_to`

`tail` parameter, which indicates the number of requests that have been processed within the last five seconds.

f) Check the tengine service.

A. Run the `docker logs --tail 50 tengine` command to check for exception logs.

B. Run the `tail -50 logs/nginx/error.log` command to check for exception logs.

g) Check the mario service.

A. Run the `docker logs --tail 50 mario` command to check for exception logs.

B. Run the `tail -50 logs/mario/mario.log` command to check for exception logs.

C. Run the `curl 127.0.0.1:3335/api/v1/state` command to check whether the service responds normally and whether the real-time request processing data is normal. For example, check whether the `num_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.

2.11.7 Routine operations and maintenance of Sensitive Data Discovery and Protection

2.11.7.1 Check the service status

2.11.7.1.1 Basic inspection

During the basic inspection of Sensitive Data Discovery and Protection (SDDP), check whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Choose **Operations > Project Operations**.
3. In the **Fuzzy Search** field, enter `yundun-sddp`.
4. Click **Details** in the **Actions** column of the `yundun-sddp` project to go to the **Cluster Operations** page.

5. In the cluster list, click the cluster name that starts with SddpCluster.
6. In the Service Instances section of the Cluster Dashboard page, check whether the yundun-sddp service instance is in the final status.

2.11.7.1.2 Advanced inspection: Check the status of the SddpService service

This topic describes how to check the running status of the SddpService service.

Procedure

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

2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.

- a) Choose **Operations > Project Operations**.
- b) In the **Fuzzy Search** field, enter `yundun-sddp`. Click **Details** in the **Actions** column of the `yundun-sddp` project to go to the **Cluster Operations** page.
- c) In the cluster list, click the cluster name that starts with `SddpCluster`.
- d) In the **Service Instances** section, find `yundun-sddp` and click **Details** in the **Actions** column to go to the **Service Instance Information Dashboard** page.

Service Instances					
Service Instance	Final Status	Expected Server Roles	Server Roles In Final ...	Server Roles Going O...	Actions
hids-client	True	1	1	0	Actions ▾ Details
os	True	--	--	--	Actions ▾ Details
tianji	True	1	1	0	Actions ▾ Details
tianji-dockerdaemon	True	1	1	0	Actions ▾ Details
yundun-sddp	True	9	9	0	Actions ▾ Details

- e) In the **Server Role List** section, find `SddpService#` and click **Details** in the **Actions** column to go to the **Server Role Dashboard** page.

Server Role List							
Server Role	Current Status	Expected Machi...	Machines In Fin...	Machines Goin...	Rolling Task St...	Time Used	Actions
SddpAlgorithm#	In Final Status	1	1	0	no rolling		Details
SddpData#	In Final Status	2	2	0	no rolling		Details
SddpDatamask#	In Final Status	2	2	0	no rolling		Details
SddpDbInit#	In Final Status	1	1	0	no rolling		Details
SddpLog#	In Final Status	2	2	0	no rolling		Details
SddpPrivilege#	In Final Status	2	2	0	no rolling		Details
SddpRuleEngine#	In Final Status	2	2	0	no rolling		Details
SddpService#	In Final Status	2	2	0	no rolling		Details
ServiceTest#	In Final Status	1	1	0	no rolling		Details

- f) In the **Machine Information** section, click **Terminal** in the **Actions** column to log on to the two physical servers of SDDP, respectively.

Machine Information									
Machi...	IP	Mac...	Mac...	Serv...	Serv...	Curr...	Targ...	Error...	Actions
a56g101...	10.10.10.10	good		good P...		2fb869ef...	2fb869ef...		Terminal Restart Details Machine System View Machine Operation
a56h1116...	10.10.10.10	good		good P...		2fb869ef...	2fb869ef...		Terminal Restart Details Machine System View Machine Operation

3. Log on to two Docker containers of the SddpService service, respectively.

Run the `sudo docker exec -it $(sudo docker ps | grep SddpService | awk '{print $1}')` bash command.

4. Check the process status of the SddpService service.

Run the `ps aux | grep java | grep yundun-sddp-service` command. If any record is returned, the service is normal.

```
#ps aux | grep java | grep yundun-sddp-service
root      162  0.1 30.7 7224188 2579604 ?        S1   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.

```
#curl 127.0.0.1:7001/checkpreload.htm
"success"
```

6. View related logs.

- View common logs in the `/home/admin/yundun-sddp-service/logs/common-log.log` file.
- View application logs in the `/home/admin/yundun-sddp-service/logs/application.log` file.
- View front-end request logs in the `/home/admin/yundun-sddp-service/logs/common-request.log` file.
- View system logs in the `/home/admin/yundun-sddp-service/logs/service-stdout.log` file.

2.11.7.1.3 Advanced inspection: Check the status of the SddpData service

This topic describes how to check the running status of the SddpData service.

Procedure

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

2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.
 - a) Choose **Operations > Project Operations**.
 - b) In the **Fuzzy Search** field, enter `yundun-sddp`. Click **Details** in the **Actions** column of the `yundun-sddp` project to go to the **Cluster Operations** page.
 - c) In the cluster list, click the cluster name that starts with `SddpCluster`.
 - d) In the **Service Instances** section, find `yundun-sddp` and click **Details** in the **Actions** column to go to the **Service Instance Information Dashboard** page.
 - e) In the **Server Role List** section, find `SddpData#` and click **Details** in the **Actions** column to go to the **Server Role Dashboard** page.
 - f) In the **Machine Information** section, click **Terminal** in the **Actions** column to log on to the two physical servers of SDDP, respectively.
3. Log on to two Docker containers of the `SddpData` service, respectively.

Run the `sudo docker exec -it $(sudo docker ps | grep SddpData | awk '{print $1}')` bash command.
4. Check the process status of the `SddpData` service.

Run the `ps aux | grep yundun-sddp-data` command. If any record is returned, the service is normal.
5. View related logs.

View logs in the `/home/admin/yundun-sddp-data/logs/sddp.log` file.

2.11.7.1.4 Advanced inspection: Check the status of the `SddpPrivilege` service

This topic describes how to check the running status of the `SddpPrivilege` service.

Procedure

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

2. Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.

- a) Choose Operations > Project Operations.
- b) In the Fuzzy Search field, enter `yundun-sddp`. Click Details in the Actions column of the `yundun-sddp` project to go to the Cluster Operations page.
- c) In the cluster list, click the cluster name that starts with `SddpCluster`.
- d) In the Service Instances section, find `yundun-sddp` and click Details in the Actions column to go to the Service Instance Information Dashboard page.
- e) In the Server Role List section, find `SddpPrivilege#` and click Details in the Actions column to go to the Server Role Dashboard page.
- f) In the Machine Information section, click Terminal in the Actions column to log on to the two physical servers of SDDP, respectively.

3. Log on to two Docker containers of the `SddpPrivilege` service, respectively.

Run the `sudo docker exec -it $(sudo docker ps | grep SddpPrivilege | awk '{print $1}')` bash command.

4. Check the process status of the `SddpPrivilege` service.

Run the `ps aux | grep java | grep yundun-sddp-privilege` command. If any record is returned, the service is normal.

5. Check the health status.

Run the `curl 127.0.0.1:7001/checkpreload.htm` command. If the response is success, the service is normal.

6. View related logs.

- View exception logs in the `/home/admin/yundun-sddp-privilege/logs/exception.log` file.
- View application logs in the `/home/admin/yundun-sddp-privilege/logs/application.log` file.
- View task logs in the `/home/admin/yundun-sddp-privilege/logs/task.log` file.
- View system logs in the `/home/admin/yundun-sddp-privilege/logs/service-stdout.log` file.

2.11.7.1.5 Advanced inspection: Check the status of the SddpLog service

This topic describes how to check the running status of the SddpLog service.

Procedure

1. *Log on to the Apsara Infrastructure Management Framework console.*
2. **Log on to the two physical servers of Sensitive Data Discovery and Protection (SDDP), respectively.**
 - a) **Choose Operations > Project Operations.**
 - b) **In the Fuzzy Search field, enter yundun-sddp. Click Details in the Actions column of the yundun-sddp project to go to the Cluster Operations page.**
 - c) **In the cluster list, click the cluster name that starts with SddpCluster.**
 - d) **In the Service Instances section, find yundun-sddp and click Details in the Actions column to go to the Service Instance Information Dashboard page.**
 - e) **In the Server Role List section, find SddpLog# and click Details in the Actions column to go to the Server Role Dashboard page.**
 - f) **In the Machine Information section, click Terminal in the Actions column to log on to the two physical servers of SDDP, respectively.**
3. **Log on to two Docker containers of the SddpLog service, respectively.**

Run the `sudo docker exec -it $(sudo docker ps | grep SddpLog | awk '{print $1}')` bash command.
4. **Check the process status of the SddpLog service.**

Run the `ps aux | grep java | grep yundun-sddp-log`. If any record is returned, the service is normal.
5. **Check the health status.**

Run the `curl 127.0.0.1:7001/checkpreload.htm` command. If the response is success, the service is normal.

6. View related logs.

- **View exception logs in the** `/home/admin/yundun-sddp-log/logs/exception.log` **file.**
- **View application logs in the** `/home/admin/yundun-sddp-log/logs/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.**

2.11.7.2 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.

2.11.8 Routine operations and maintenance of Apsara Stack Security Center

2.11.8.1 Check service status

2.11.8.1.1 Basic inspection

During the basic inspection of Apsara Stack Security Center, check whether the service has reached the final status.

Procedure

1. [Log on to the Apsara Infrastructure Management Framework console.](#)
2. Choose **Operations > Project Operations**. Enter `yundun-advance`, and click **Details** to go to the **Cluster Operations** page.
3. Select **BasicCluster**.
4. Check whether `yundun-secureconsole` has reached the final status in **Service Instances List**.

2.11.8.1.2 Advanced inspection

Check the running status of Apsara Stack Security Center.

Context

To check the running status of Apsara Stack Security Center, follow the following steps:

Procedure

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

2. Log on to two physical machines, respectively.
 - a) Choose **Operations > Project Operations**.
 - b) Enter `yundun-advance`, and click **Details** to go to the **Cluster Operations** page.
 - c) Select **BasicCluster**.
 - d) Select `yundun-secureconsole` from **Service Instances List**, and click **Details** to go to the **Service Instance Dashboard** page.
 - e) Select `SecureConsoleApp#` from **Service Role List**, and click **Details** to go to the **Service Role Dashboard** page.
 - f) View **Server Information**, and use **TerminalService** to log on to two physical machines, respectively.

3. Log on to two secure-console Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep secureconsole | awk '{print $1}') bash.
```

4. Check the console progress status.

Run `ps aux |grep console`. If any record is returned, the console progress is normal.

5. Check the health status.

Run `curl 127.0.0.1:3014/check.htm`. If **OK** is returned, the service is normal.

6. View related logs.

- View the Tomcat logs in `/home/admin/console/logs/jboss_stdout.log`.

2.11.8.2 Restart the secure-console service

Context

To restart the secure-console service when an error occurs, follow the following steps:

Procedure

1. Run the `ssh server IP address` command to log on to the server that hosts the secure-console service.
2. Run the following command to find the image ID of the secure-console service:

```
sudo docker ps -a |grep console
```

3. Run the following command to go to the Docker container:

```
docker exec -it [imageId] /bin/bash
```

4. Run the following command to locate the Java process:

```
ps aux |grep console
```

5. Run the following command to stop the current process:

```
kill -9 process
```

6. Run the following command to restart the process:

```
sudo -u admin /home/admin/console/bin/jbossctl restart
```

7. Run the following command to check whether the process has been successfully restarted:

```
curl 127.0.0.1:7001/check.htm
```

2.11.9 Routine operations and maintenance of secure-service

2.11.9.1 Check the service status

2.11.9.1.1 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.**

2.11.9.1.2 Advanced inspection: Check the secure-service status

This topic describes how to check the secure-service running status.

Procedure

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

2. Log on to two physical machines, respectively.

- a) **Choose Operations > Project Operations.**
- b) **Enter yundun-advance, and click Details to go to the Cluster Operations page.**
- c) **Select BasicCluster.**
- d) **Select yundun-secureservice from Service Instances List, and click Details to go to the Service Instance Dashboard page.**
- e) **Select SecureServiceApp# from Service Role List, and click Details to go to the Service Role Dashboard page.**
- f) **View Server Information, and click Terminal to log on to two physical machines, respectively.**

3. Log on to two secure-service Docker containers, respectively.

```
Run sudo docker exec -it $(sudo docker ps | grep secureservice | awk '{print $1}') bash.
```

4. Check the secure-service process status.

Run `ps aux |grep secure-service`. If any record is returned, the secure-service process is normal.

5. Check the health status.

Run `curl 127.0.0.1:3010`. If OK is returned, the service is normal.

6. Run the following command to go to the Docker container:

```
sudo docker exec -it [imageId] /bin/bash
```

7. View related logs.

- **View the Server Guard logs in** `/home/admin/secure-service/logs/aegis-info.log`.
- **View the error logs in** `/home/admin/secure-service/logs/Error`.
- **View the vulnerability analysis and scanning logs in** `/home/admin/secure-service/logs/leakage-info.log`.
- **View the cloud intelligence logs in** `/home/admin/secure-service/logs/threat-info.log`.
- **View the web attack logs in** `/home/admin/secure-service/logs/web-info.log`.
-

2.11.9.1.3 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/dolphin-message-producer.log`.



Note:

Currently, the Dolphin service sends messages only to TDS.

2.11.9.1.4 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`.

2.11.9.2 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
```

2.12 Key Management Service (KMS)

2.12.1 Operations and maintenance of KMS components

2.12.1.1 Overview

KMS is deployed and managed from the Apsara Infrastructure Management Framework console. On the Machine Operations page of the Apsara Infrastructure Management Framework console, you can access a host where KMS is deployed.

2.12.1.2 KMS_HOST

Determine whether the service role functions properly

- 1. In the Apsara Infrastructure Management Framework console, check whether the KMS_HOST service role has been deployed.**

Follow these steps:

- a. Log on to the Apsara Infrastructure Management Framework console.**
- b. In the top navigation bar, choose Tasks > Deployment Summary. The Deployment Summary page is displayed.**
- c. Click Deployment Details.**
- d. On the Deployment Details page, locate kms.**
- e. In the top navigation bar, choose TasksDeployment Summary. On the Deployment Summary page that appears, click Deployment Details to view the**

deployment status of the KMS_HOST service role, as shown in *Figure 2-26: Check whether the KMS_HOST service role has been deployed.*

If a tick next to KMS_HOST# is green, the KMS_HOST service role has been deployed.

Figure 2-26: Check whether the KMS_HOST service role has been deployed

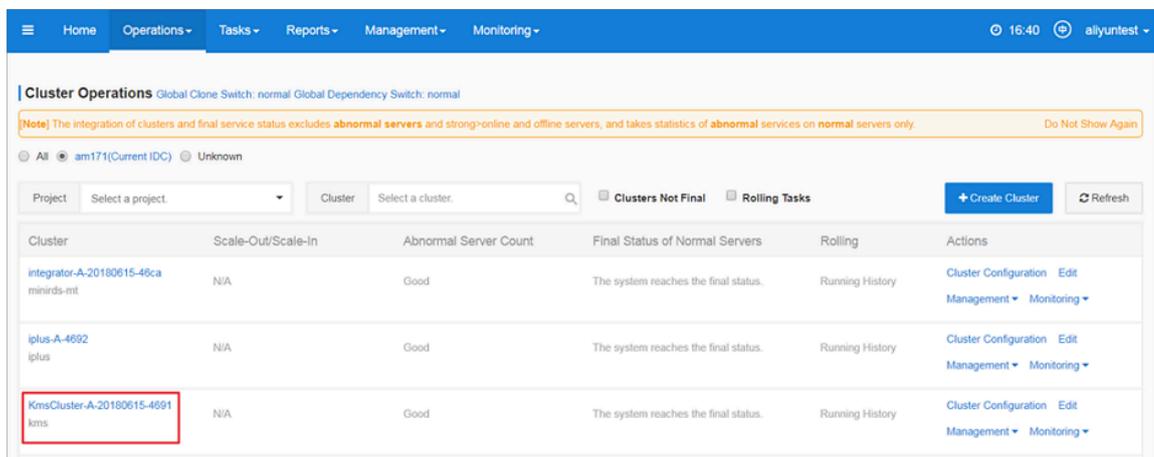
kms		Final	4 Days 13 Hours	Cluster: 1 / 1	Service: 5 / 5	Role: 11 / 11	Total: 6
lark	Final	1 Day		...	hids-client	✓	Etcd# ✓
middleWare-histore	Final	1 Day		...	kms	✓	EtcdDecider# ✓
middleWare-queqiao	Final	12 Hours 54 Minutes		...	os	✓	HSA# ✓
middleWare-redis	Final	4 Days 14 Hours		...	tianji	✓	KmsHost# ✓
middleWare-staragent	Final	1 Day		...	tianji-dockerdae...	✓	KmsInit# ✓
middleWareAll	Final	1 Day		...			Rotator# ✓
middleware-dnccs	Final	1 Day 23 Hours		...			ServerroleMonitor# ✓
							ServiceTest# ✓

2. Obtain the IP addresses of the hosts where the KMS_HOST service role has been deployed.

Follow these steps:

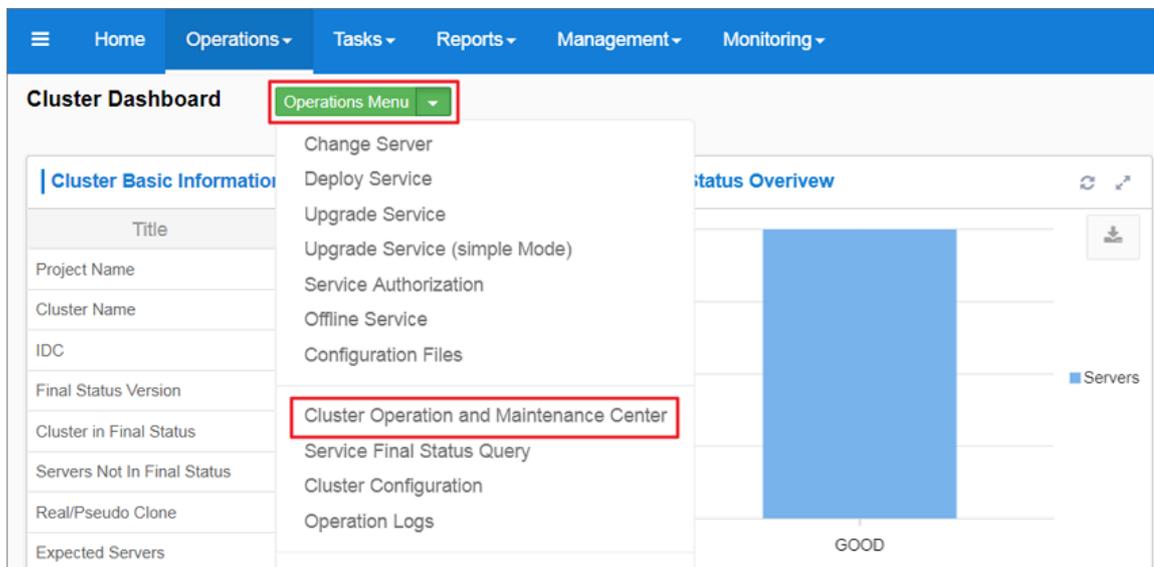
- a. Log on to the Apsara Infrastructure Management Framework console.**
- b. In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page that appears, search for the corresponding cluster, as shown in *Figure 2-27: Search for clusters*.**

Figure 2-27: Search for clusters



- c. Click a specified cluster URL to go to the Cluster Dashboard page.**
- d. On the Cluster Dashboard page, choose Operations Menu > Cluster Operation and Maintenance Center, as shown in *Figure 2-28: Cluster Dashboard*.**

Figure 2-28: Cluster Dashboard



- e. On the Cluster Operation and Maintenance Center page, view and obtain the IP addresses of all hosts where the KMS_HOST service role has been deployed, as shown in *Figure 2-29: View IP addresses of hosts*.

Figure 2-29: View IP addresses of hosts

Cluster Operations > Cluster Operation and Maintenance Center

Cluster Operation and Maintenance Center (Cluster: KmsCluster-A-20180615-4691)

SR not in Final Status: *N/A* Running Tasks: No rolling is availab... Head Version Commit Time: 20/07/18, 02:31:31 Head Version Analysis Status: **done**

Service: **kms** Service Role: **HSA#**

Total Servers: **3** Expected Ser...: 3 Scale-In Scale-Out: Server: s-in: 0 s-out: 0 SR: s-in: 0 s-out: 0 Abnormal Server: **0** Ping Failed: 0 No Heartbeat: 0 Status Error: 0 Abnormal Service: **0** TJ-Client: 0 Other SRs: 0

Server List

Server Search: Supports multiple server search. [Reset]

Server	Final SR Status	SR Running Status	Action	Action Status	Monitoring Statistics	Actions
a56e07014.cloud.e07.am171 10.12.2.21	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 8	Terminal Approval Action Restart Service Role
a56e11114.cloud.e12.am171 10.12.4.13	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 8	Terminal Approval Action Restart Service Role
a56f11114.cloud.f12.am171 10.12.6.20	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 8	Terminal Approval Action Restart Service Role

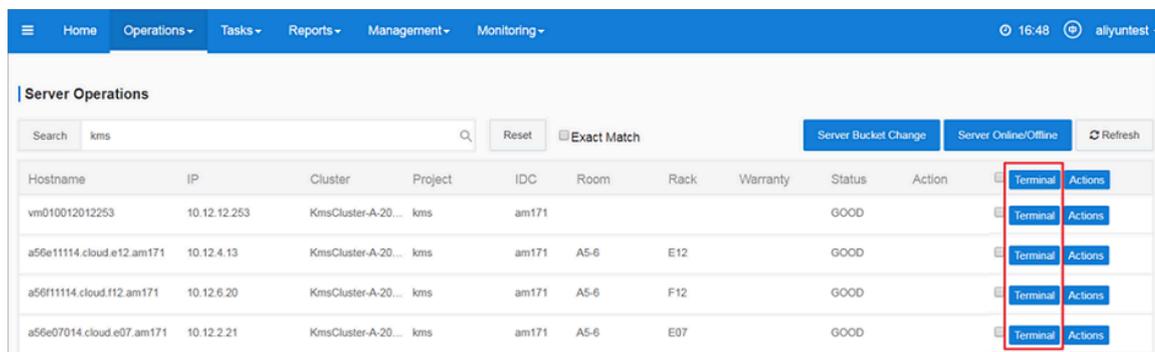
[Batch Terminal] Items per Page: 10 1

3. Access the KMS server and run the `curl http://ip:5555/status.html` command to check whether success is returned.

Follow these steps:

- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose Operations > Machine Operations. On the Machine Operations page that appears, locate a relevant host, as shown in [Figure 2-30: Search for hosts](#).

Figure 2-30: Search for hosts



Hostname	IP	Cluster	Project	IDC	Room	Rack	Warranty	Status	Action
vm010012012253	10.12.12.253	KmsCluster-A-20...	kms	am171				GOOD	Terminal Actions
a56e11114.cloud.e12.am171	10.12.4.13	KmsCluster-A-20...	kms	am171	A5-6	E12		GOOD	Terminal Actions
a56f11114.cloud.f12.am171	10.12.6.20	KmsCluster-A-20...	kms	am171	A5-6	F12		GOOD	Terminal Actions
a56e07014.cloud.e07.am171	10.12.2.21	KmsCluster-A-20...	kms	am171	A5-6	E07		GOOD	Terminal Actions

- c. Select a host and click Terminal to log on to the host through a terminal session.
- d. Run the `curl http://ip:5555/status.html` command to check whether success is returned, as shown in [Figure 2-31: Enter a command](#). Verify all hosts where the KMS_HOST service role has been deployed based on the previous procedure.

IP indicates the IP addresses of the hosts that are obtained in the previous step

.

Figure 2-31: Enter a command



```
$ curl http://[redacted]:5555/status.html
success
```

Locate and determine exceptions

1. View log entries in `/cloud/log/kms/KmsHost#/kms_host`.
2. Check whether the KMS_HOST service role functions properly. If the KMS_HOST service role exits soon after it is started, check `debug.log` to locate the cause.

3. If the `KMS_HOST` service role runs properly with faulty functions, view `status.log` to locate the cause.

Possible exceptions and errors

- `xxx selfCheck error`



Note:

`xxx` indicates a dependent service.

1. Check whether the corresponding dependency configurations are correct. You can use `debug.log` to locate the configurations.
2. Check whether `xxx` runs properly.

- `exit code 1`

Locate the cause of an unexpected exit based on `debug.log`.

2.12.1.3 HSA

Check whether the server role is working normally

1. In the Apsara Infrastructure Management Framework console, check whether `HSA#` has reached the final state.

Follow these steps:

- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose `Tasks > Deployment Summary`. The `Deployment Summary` page is displayed.
- c. Click `Deployment Details`.
- d. On the `Deployment Details` page, locate `kms`.
- e. In the `Deployment Progress` column corresponding to `kms`, move the pointer to the right of `Role`. Click `Details` that appears. Check whether `HSA#` has

reached the final state, as shown in *Figure 2-32: Check whether HSA# has reached the final state.*

If a green tick appears for HSA#, HSA# has reached the final state.

Figure 2-32: Check whether HSA# has reached the final state

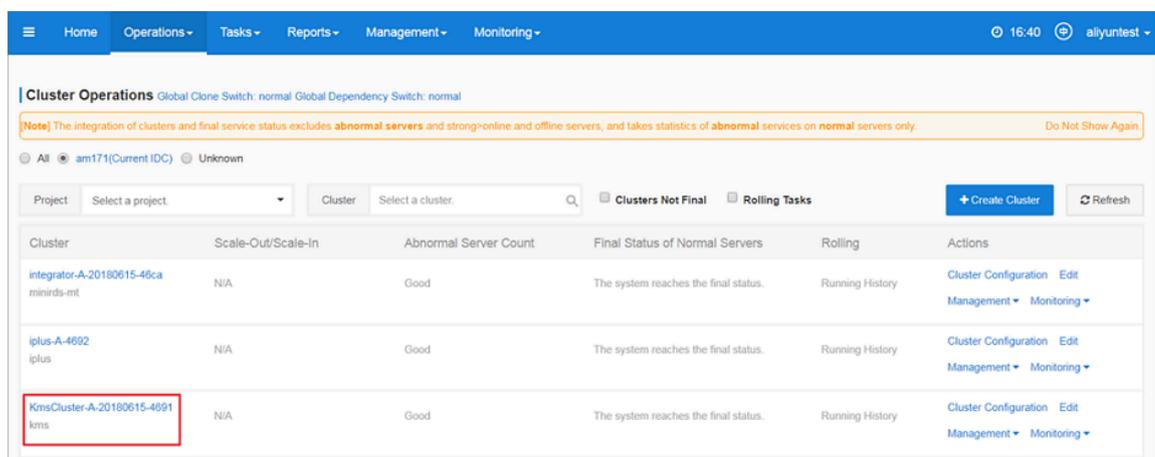
kms		Final	4 Days 13 Hours	Cluster: 1 / 1	Service: 5 / 5	Role: 11 / 11	Total: 6
lark	Final	1 Day		⚙️ KmsCluster-A-2...	⚙️ hids-client	⚙️ Etc#	⚙️
middleWare-histore	Final	1 Day			⚙️ kms	⚙️ EtcDecider#	⚙️
middleWare-queqiao	Final	12 Hours 54 Minutes			⚙️ os	⚙️ HSA#	⚙️
middleWare-redis	Final	4 Days 14 Hours			⚙️ tianji	⚙️ KmsHost#	⚙️
middleWare-staragent	Final	1 Day			⚙️ tianji-dockerdae...	⚙️ KmsInit#	⚙️
middleWareAll	Final	1 Day				⚙️ Rotator#	⚙️
middleware-dnscs	Final	1 Day 23 Hours				⚙️ ServerroleMonitor#	⚙️
						⚙️ ServiceTest#	⚙️

2. Obtain the IP addresses of the machines that are deployed with the KmsHost# service.

Follow these steps:

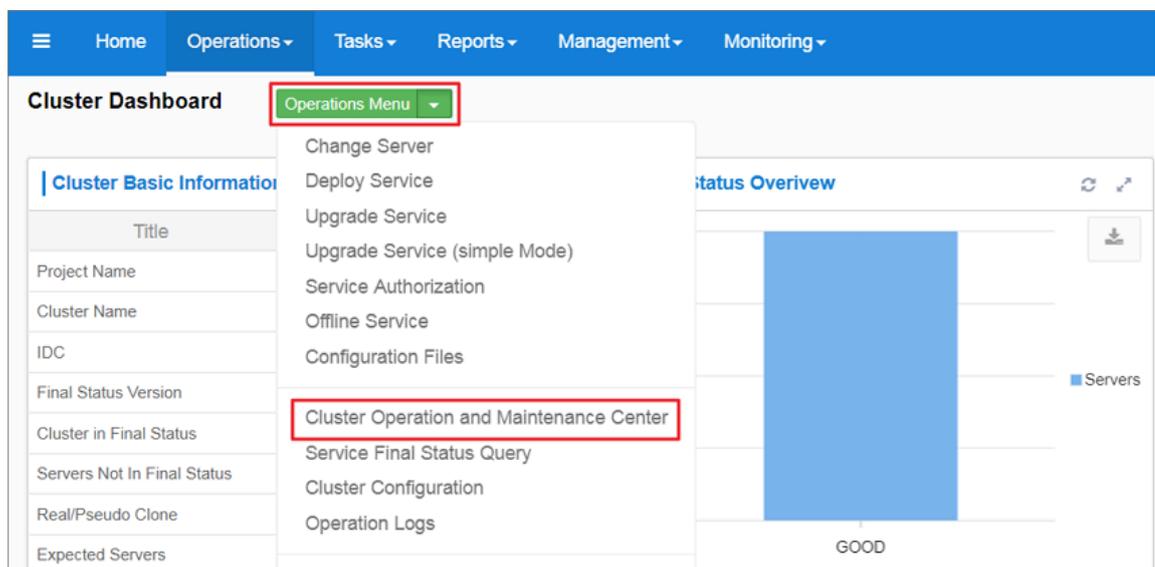
- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose **Operations > Cluster Operations**. On the **Cluster Operations** page that appears, search for the corresponding cluster, as shown in *Figure 2-33: Search for clusters*.

Figure 2-33: Search for clusters



- c. Click a specified cluster URL to go to the **Cluster Dashboard** page.
- d. On the **Cluster Dashboard** page, choose **Operations Menu > Cluster Operation and Maintenance Center**, as shown in *Figure 2-34: Cluster Dashboard*.

Figure 2-34: Cluster Dashboard



- e. On the Cluster Operation and Maintenance Center page, view and obtain the IP addresses of all machines that are deployed with the HSA service, as shown in *Figure 2-35: Obtain the IP addresses of the machines that are deployed with the HSA service.*

Figure 2-35: Obtain the IP addresses of the machines that are deployed with the HSA service

The screenshot displays the 'Cluster Operation and Maintenance Center' for cluster 'KmsCluster-A-20180615-4691'. The service is 'kms' and the role is 'HSA#'. The summary shows 3 total servers, 0 abnormal servers, and 0 abnormal services. The server list table is as follows:

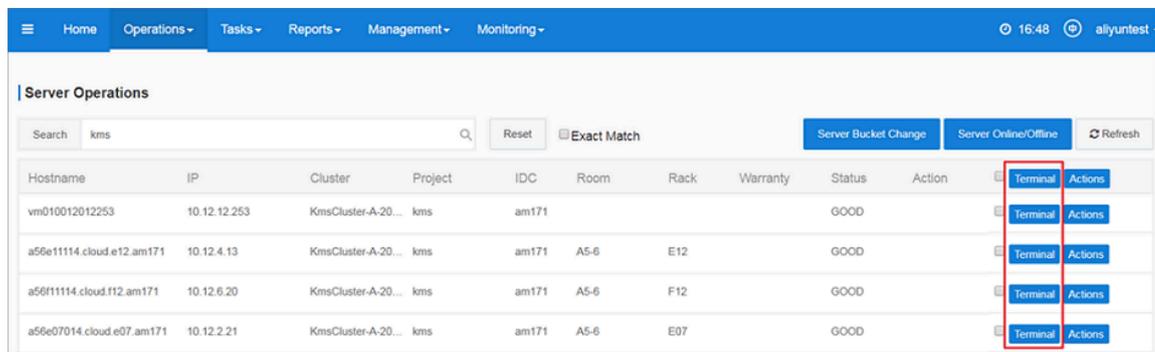
Server	Final SR Status	SR Running Status	Action	Action Status	Monitoring Statistics	Actions
a56e07014.cloud.e07.am171 10.12.2.21	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 8	Terminal, Approval Action, Restart Service Role
a56e11114.cloud.e12.am171 10.12.4.13	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 8	Terminal, Approval Action, Restart Service Role
a58f11114.cloud.f12.am171 10.12.6.20	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 8	Terminal, Approval Action, Restart Service Role

3. Log on to the KMS server and run the `curl http://ip:8081/status.html` command. Check whether success is returned. If yes, the server role is working normally.

Follow these steps:

- Log on to the Apsara Infrastructure Management Framework console.
- In the top navigation bar, choose **Operations > Machine Operations**. On the Machine Operations page that appears, locate a relevant host, as shown in [Figure 2-36: Search for hosts](#).

Figure 2-36: Search for hosts



Hostname	IP	Cluster	Project	IDC	Room	Rack	Warranty	Status	Action
vm010012012253	10.12.12.253	KmsCluster-A-20...	kms	am171				GOOD	Terminal Actions
a56e11114.cloud.e12.am171	10.12.4.13	KmsCluster-A-20...	kms	am171	A5-6	E12		GOOD	Terminal Actions
a56f11114.cloud.f12.am171	10.12.6.20	KmsCluster-A-20...	kms	am171	A5-6	F12		GOOD	Terminal Actions
a56e07014.cloud.e07.am171	10.12.2.21	KmsCluster-A-20...	kms	am171	A5-6	E07		GOOD	Terminal Actions

- Select a host and click **Terminal** to log on to the host through a terminal session.
- Run the `curl http://ip:8081/status.html` command to verify whether the machine is deployed with the HSA service. If the server role is working normally, success is returned for each machine, as shown in [Figure 2-37: Enter a command](#).

In the command, `ip` indicates the IP address obtained in the previous step. It is the IP address of the machine deployed with the HSA service.

Figure 2-37: Enter a command

```
$curl http://[redacted]:8081/status.html
success
```

Troubleshooting

- View the relevant log files in `/cloud/log/kms/HSA#/hsa`.

2. Check whether the hsa application works normally. If it exits soon after startup, view `debug.log` to identify the issue.
3. If the hsa application works normally but does not function properly, view `status.log` to identify the issue.

Possible errors and exceptions

Error: `exit code 1`

View `debug.log` to identify the cause of the exceptional exit.

Common possible causes include:

- `etcd` has not been started normally.
- `etcd` has been started normally but there is no valid data.



Note:

During disaster recovery, synchronization errors in the backup cluster may cause this error.

2.12.1.4 etcd

Check whether the server role is working normally

In the Apsara Infrastructure Management Framework console, check whether `Etcd` # and `EtcdDecider`# have reached the final state.

Follow these steps:

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the top navigation bar, choose `Tasks > Deployment Summary`. The `Deployment Summary` page is displayed.
3. Click `Deployment Details`.
4. On the `Deployment Details` page, locate `kms`.
5. In the `Deployment Progress` column corresponding to `kms`, move the pointer to the right of `Role`. Click `Details` that appears. Check whether `Etcd`# and

EtcdDecider# have reached the final state, as shown in *Figure 2-38: Check whether Etcd# and EtcdDecider# have reached the final state.*

If a green tick appears for both **Etcd#** and **EtcdDecider#**, they have reached the final state.

Figure 2-38: Check whether Etcd# and EtcdDecider# have reached the final state

Service Name	Status	Duration	Cluster	Service	Role	Total
kms	Final	4 Days 13 Hours	1 / 1	5 / 5	11 / 11	6
lark	Final	1 Day				
middleWare-histore	Final	1 Day				
middleWare-queqiao	Final	12 Hours 54 Minutes				
middleWare-redis	Final	4 Days 14 Hours				
middleWare-staragent	Final	1 Day				
middleWareAll	Final	1 Day				
middleware-dnscs	Final	1 Day 23 Hours				

Troubleshooting

- etcd exceptions are complex. The log file in `/cloud/log/kms/Etcd#/etcd` only provides partial information.**
- To find out the rest information about the exceptions, view the log file in `/cloud/log/kms/EtcdDecider#/decider`.**

Possible errors and exceptions

The possible errors and exceptions are as follows:

- Errors occur in the startup parameters of etcd for some special reasons.**

Retain the log files of etcd and contact Alibaba Cloud technical support personnel to find the reasons.

Quick solution: Find correct startup parameters in `debug.log` of etcd and start etcd manually.

- Errors in EtcdDecider during service upgrades cause errors in etcd.**

Typically, this occurs when there is a rolling task. You can analyze the issue and identify the cause based on the `debug.log` file of EtcdDecider.

- The data directory of etcd is missing and etcd cannot start.

Solution: Use the Apsara Infrastructure Management Framework console to scale in the abnormal etcd node, and then scale out the node to its original server role group.

2.12.1.5 Rotator

2.12.1.5.1 Primary IDC

The status of the Rotator service role is special. Even if the Apsara Infrastructure Management Framework console shows that the service role has been deployed, the Rotator service role is not necessarily working properly.

Rotator service role exceptions do not have any adverse impact on the API logic of KMS.

Typically, you must use log entries to locate the cause of a fault for the Rotator service role only after unexpected results are found. For example, the data on RDS does not match expectations.

Determine whether the Rotator service role is enabled in the primary IDC mode

View `current idc master` in `/cloud/log/kms/Rotator#/rotator/debug.log`, as shown in [Figure 2-39: View the IDC mode](#). Determine whether the Rotator service role (in the primary IDC) is enabled in the primary IDC mode.

Figure 2-39: View the IDC mode

```
[2017-10-16 13:07:50.458588] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] PKIVersion:pssl
[2017-10-16 13:07:50.497312] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] current idc master: true
[2017-10-16 13:07:50.553460] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] CurrentClients:map[a27d05007.ccloud.d05.ew9-5:0xc4206d6720 a27d08007.ccloud.d08.ew9-5:0xc420647da0 a27d11007.ccloud.d11.ew9-5:0xc4206d7980]
```

If the `current idc master` value indicates `true`, the Rotator service role is enabled in the primary IDC mode. If the `current idc master` value indicates `false`, the Rotator service role is enabled in the secondary IDC mode.

Check whether the Rotator service role is in the working state

The Rotator service role of the primary IDC is deployed to all nodes in the distributed lock mode. In this mode, only one node is in the working state and the other nodes are all in the standby state.

View `/cloud/log/kms/Rotator#/rotator/status.log` to determine whether the Rotator service role is in the working state.

As shown in [Figure 2-40: Working state](#) and [Figure 2-41: Standby state](#):

- **ExecuteWorker:** The node is in the working state.
- **TryLock:** The node is in the standby state.

Figure 2-40: Working state

```
[2017-10-23 16:51:51.554310] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d05007.cloud.d05.ew9-5 RotatorState:ExecuteWorker
[2017-10-23 16:52:51.554415] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d05007.cloud.d05.ew9-5 RotatorState:ExecuteWorker
```

Figure 2-41: Standby state

```
11007.cloud.d11.ew9-5 RotatorState:TryLock
[2017-10-17 18:35:20.618575] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d11007.cloud.d11.ew9-5 RotatorState:TryLock
[2017-10-17 18:36:11.867967] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d11007.cloud.d11.ew9-5 RotatorState:TryLock
[2017-10-17 18:36:20.620963] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] module:Rotator host:a27d11007.cloud.d11.ew9-5 RotatorState:TryLock
```

Possible exceptions and errors

- **Abnormal RDS database access.** The statistic collection and key deletion tasks cannot be run properly.
- **Abnormal HSA service role.** The key update task cannot be run properly.
- **Abnormal Log Service.** The metering task cannot be run properly.
- **Abnormal Etcd service role.** The distributed lock is unavailable, and tasks cannot be run.
- **If one of the tasks on the Rotator service role is abnormal, the Rotator service role may be unable to be deployed in the Apsara Infrastructure Management Framework console.**

2.12.1.5.2 Secondary IDC

The Rotator service role of the secondary IDC is deployed to all nodes. Each node is in the working state. The work scope of each node is idempotent to those of other nodes within a certain time range.

Determine whether the Rotator service role is enabled in the secondary IDC mode

View `current idc master` in `/cloud/log/kms/Rotator#/rotator/debug.log`, as shown in [Figure 2-42: View the IDC mode](#). Determine whether the Rotator service role (in the secondary IDC) is enabled in the secondary IDC mode.

Figure 2-42: View the IDC mode

```
[2017-10-21 16:34:34.412535] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] PKIVersion:pssl
[2017-10-21 16:34:34.446620] [INFO] [logger.(*LoggerWrapper).Infof] [logwrapper.go:37] current idc master: false
```

If the `current idc master` value indicates `false`, the Rotator service role is enabled in the secondary IDC mode. If the `current idc master` value indicates `true`, the Rotator service role is enabled in the primary IDC mode.

Possible exceptions and errors

- **Abnormal primary IDC networks.** The Etcd service role of the primary IDC is inaccessible.
- **Abnormal Etcd service role of the primary IDC.** The Etcd service role of the primary IDC is inaccessible.
- **Abnormal Etcd service role of the secondary IDC.** Data write fails.
- **Incorrect Etcd service role information of the primary IDC.** An error occurs during data synchronization.



Notice:

Rotator service role exceptions of the secondary IDC severely affect the KMS in the primary IDC. Handle this exception immediately.

2.12.2 Log analysis

2.12.2.1 Overview

Logtail is a log collection client provided by Log Service to facilitate your access to logs. After installing Logtail on a host that has KMS deployed, you can monitor a specified log. The newly written log entries are automatically uploaded to a specified log library.

Logtail is used to transmit the logs of KMS to Log Service. Then the portal or API of Log Service analyzes the logs. If Log Service has no portals, you have to log on to the hosts that have KMS deployed individually and check the hosts one by one.

2.12.2.2 Request IDs

After sending a request to KMS, you will receive a response from KMS. The response contains a request ID.

Request IDs are used in the following scenarios:

- Go to `/cloud/log/kms/KmsHost#/kms_host/audit.log` to view the KMS audit log.

You can use the `request_id` value to view the audit log information of the current access.

- For log entries whose `expected_code` values are not 200, you can view error information during debugging based on the request ID.

Path to the local log: `/cloud/log/kms/KmsHost#/kms_host/debug.log`



Note:

`/cloud/log/kms/KmsHost#/kms_host/debug.log` and `audit.log` are stored in the same host.

- If you need all details of a request, you can view detailed information in the trace log.

Path to the local log: `/cloud/log/kms/KmsHost#/kms_host/debug.log`



Note:

`/cloud/log/kms/KmsHost#/kms_host/debug.log` and `audit.log` are stored in the same host.

- You can use the request ID to associate cryptography-relevant APIs with the trace log of HSA.

Path to the local log: `/cloud/log/kms/HSA#/hsa/trace.log`



Note:

`/cloud/log/kms/KmsHost#/kms_host/trace.log` and `audit.log` may be stored in different hosts.

- You can also retrieve logs based on other information.

You can retrieve audit logs of KMS based on other information. However, you still need the request ID to associate the audit logs with other logs.

2.12.2.3 Common KMS errors

2.12.2.3.1 Overview

KMS has two HTTP status codes in `audit.log`: `expected_code` and `status_code`.

Typically, the `expected_code` and `status_code` of an error are the same. (`expected_code = status_code`). However, there are exceptions.

`status_code` is the HTTP status code that is actually returned to a user.

2.12.2.3.2 Error codes 4xx

Error codes 4xx indicate expected errors in KMS. For example, error code 403 indicates a user authentication request failure and error code 400 indicates that the parameters entered by a user are incorrect.

You can use the request ID to view detailed error information during debugging.

2.12.2.3.3 Error code 500

Typically, if the status code of an error is 500, the `expected_code` of the error is also 500.

Errors of this type are not expected by KMS. Typically, they are severe errors and must be fixed immediately.

A dependent service probably encounters unexpected errors. We recommend that you contact Customer Services to troubleshoot the problem.

You can use the request ID to view detailed error information during debugging.

2.12.2.3.4 Error code 503

Error code 503 occurs in the following scenarios:

- The `expected_code` is not 503 but the `status_code` is 503.

Possible causes:

- The client-side user interrupted the connection in advance.
- The client has timed out because the response of the KMS server is too slow.

You can use the request ID in the trace log to determine whether the server has timed out and identify the modules that have timed out.

- `expected_code=status_code=503`

An expected error occurred in a dependent service of the KMS. This is caused by an unstable dependent service.

You can use the request ID to view detailed error information during debugging. We recommend that you contact Customer Services to troubleshoot the problem.

2.12.2.3.5 Dependent service degradation

KMS caches dependent service data to local memory. If a dependent service is unavailable, KMS uses the obsolete data cached to the local device to continue providing services.

In this scenario, the status code in the audit log of KMS is 200, but an additional debug log entry is generated.

In this scenario, some users can access KMS (data cached), but other users may encounter a 503 error (data not cached).

2.12.3 View and process internal data

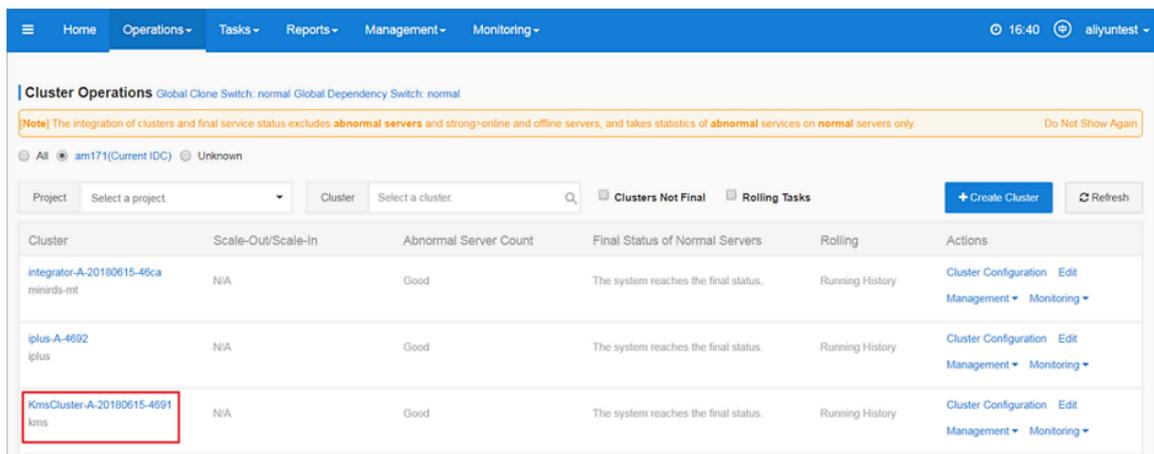
View updates to CMKs

1. Log on to the kmsdata database from an Apsara Stack server that has a MySQL client installed.

To obtain connection information of the kmsdata database, perform the following operations:

- a. Log on to the Apsara Infrastructure Management Framework console.
- b. In the top navigation bar, choose Operations > Cluster Operations. On the Cluster Operations page that appears, search for the corresponding cluster, as shown in *Figure 2-43: Search for clusters*.

Figure 2-43: Search for clusters



- c. Click a specified cluster URL to go to the Cluster Dashboard page.
- d. Scroll down the page and locate Cluster Resource.
- e. In Cluster Resource, locate the kmsdata database, as shown in *Figure 2-44: Cluster resources*.

Figure 2-44: Cluster resources

Service	serverr...	app	name	type	status	error_...	param...	result	res	reproc...	reproc...	reproc...	refer_v...
kms	kms.KmsInit#	db_init	kmsdata	db	done		{"minirds_p...	{"passwd": "...	ddec793d5...				[{"#587c17...
kms	kms.KmsInit#	db_init	kms_servic...	accesskey	done		{"name": "k...	{"name": "k...	610f802a4...				[{"#587c17...
kms	kms.KmsInit#	db_init	kms-internet	vip	done		{"check_ty...	{"nc_list": "...	4f59db9f05...				[{"#587c17...
kms	kms.KmsInit#	db_init	kms-intranet	vip	done				3e4761db5...	done			[{"#587c17...
kms	kms.KmsInit#	db_init	kms-intranet	dns	done		{"domain": "...	{"ip": "[{"10...	4ec409337...				[{"#587c17...
kms	kms.KmsInit#	db_init	kms-internet	dns	done		{"domain": "...	{"ip": "[{"42...	3fa990dcde...				[{"#587c17...

f. In the **Result** column corresponding to the cluster, right-click and choose **Show More** from the shortcut menu.

In the **Details** message that appears, you can view the connection information of the kmsdata database.

2. Enter the `select MIN(dk_version) from ekt_tbl;` SQL statement, and view the `dk_version` information, as shown in *Figure 2-45: View dk_version information*.



Note:

`dk_version` indicates the date of the current day.

Figure 2-45: View dk_version information

```
MySQL [kmsdata]> select MIN(dk_version) from ekt_tbl;
+-----+
| MIN(dk_version) |
+-----+
| DK-201710250000 |
+-----+
```

View a deleted CMK

After a CMK is deleted, it is moved to `dustbin_cmk_tbl` and `dustbin_ekt_tbl`.

Run the following SQL statements to view the information of the deleted CMK:

```
select * from dustbin_cmk_tbl limit 1;
```

```
select * from dustbin_ekt_tbl limit 1;
```


ns on Alibaba Cloud. Each console instance contains an NGINX server and a Jetty container.

Command portal

1. Log on to the Apsara Infrastructure Management Framework console.
2. In the top navigation bar, choose Operations > Service Operations.
3. Locate `sls-backend-server` and click Management in the Actions column. On the page that appears, click Service Instance.
4. Click the service instance name to go to the Service Instance Information Dashboard page. Locate `WebServer#` in Server Role List and click Details to view the corresponding machines.
5. Click Terminal in the Actions column corresponding to a machine, log on to the machine, and open the file directories to find the NGINX and Jetty services.

Commands

- Restart NGINX:

```
sudo /etc/init.d/nginx restart
```

- Restart Jetty:

```
sudo /etc/init.d/jetty restart
```

Directories

- Web application root directory: `/alidata/www/`
- Console application war directory: `/alidata/www/wwwroot/sls-console-aliyun-com/`
- Service application war directory: `/alidata/www/wwwroot/sls-service-aliyun-com/`
- Static resources directory: `/alidata/www/wwwroot/static/`

Configuration files

- NGINX: `/etc/nginx/conf.d/sls-console-aliyun-com.conf.console`
- Console: `/alidata/www/wwwroot/sls-console-aliyun-com/WEB-INF/classes/config/web.properties`
- Service: `/alidata/www/wwwroot/sls-service-aliyun-com/WEB-INF/classes/config/sls.properties`

Application logs

- **NGINX logs:** `/apsara/nginx/logs/sls_console.log`
- **Log root directory:** `/alidata/www/logs/`
- **Console application logs:** `/alidata/www/logs/java/sls/`
- **Service application logs:** `/alidata/www/logs/java/sls-service/`
- **Jetty logs:** `/usr/share/jetty/log/`

2.13.2 Troubleshooting

2.13.2.1 NGINX

Error log: `/apsara/nginx/log/error.log`

Error	Solution
Bind Address Failed	Check port listening information in <code>/etc/init.d/nginx.conf</code>.
open() ... failed	Check whether the item you want to open exists in the static resource file.

2.13.2.2 Console

Error log: `/alidata/www/logs/java/sls/error.log`

Error	Solution
SLS SDK Exception	Normal log entry. No action is required.
Create Bean Failed	Check the dubbo settings in the Console configurations.

2.13.2.3 Service

Error log: `/alidata/www/logs/java/sls-service/applog/error.log`

Error	Solution
Create Bean Failed	Check the dubbo settings in the Service configurations.
Invoke failed	Check the scmg settings in the Service configurations.

2.14 Apsara Stack DNS

2.14.1 Introduction to Apsara Stack DNS

This topic describes the features and modules of Apsara Stack DNS.

Database management system

The database management system compares versions in the baseline configurations with the versions in the database. The system ensures that the database version is accurate and up-to-date.

API system

The API system is written in Java. The system determines the business logic of all calls and manages all data and tasks.

DNS resolution system

The DNS resolution system consists of BIND and Agent. Agent receives and processes task information that is passed from the API system, parses the tasks into commands, and then sends the commands to the BIND system.

2.14.2 Maintenance

2.14.2.1 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.

2.14.2.2 Enable and disable a service

You can log on to the API server as an administrator and run the `/home/admin/gdns/bin/appctl.sh restart` command to restart the API service. We recommend that you run the command on one server at a time to ensure that another server can provide services. You can specify the `start`, `stop`, and `restart` parameters in the preceding command.

Apsara Stack DNS provides services by using anycast IP addresses. You must run the `service ospfd stop` command to disable the OSPF service before you run the `service named stop` command to disable the DNS service.

You must run the `service named start` command to enable the DNS service before you run the `service ospfd start` command to enable the OSPF service.

You can run the `/usr/local/AgentService/agent -s start` command to enable the Agent service. If you receive a message that indicates the PID file already exists, delete the `/var/dns/dns.pid` file and run the command again.

You can run the `/usr/local/AgentService/agent -s stop` command to disable the Agent service.

2.14.2.3 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.

2.14.3 DNS API

2.14.3.1 Manage the API system

You can manage the API system in the Apsara Infrastructure Management Framework console. To log on to the server in which the API system resides, choose **Operations > Server Operations in the Apsara Infrastructure Management Framework console.**

Context

To determine whether a service role is running as expected, follow these steps:

Procedure

1. In the Apsara Infrastructure Management Framework console, check whether the API is at desired state.

a) Log on to the Apsara Infrastructure Management Framework console.

b) In the top navigation bar, choose Tasks > Deployment Summary to open the Deployment Summary page.

c) Click Deployment Details.

d) On the Deployment Details page, find the dnsProduct project.

e) Find the dnsServerRole# service role, and click Details in the Deployment Progress column to check whether the service role is at desired state.

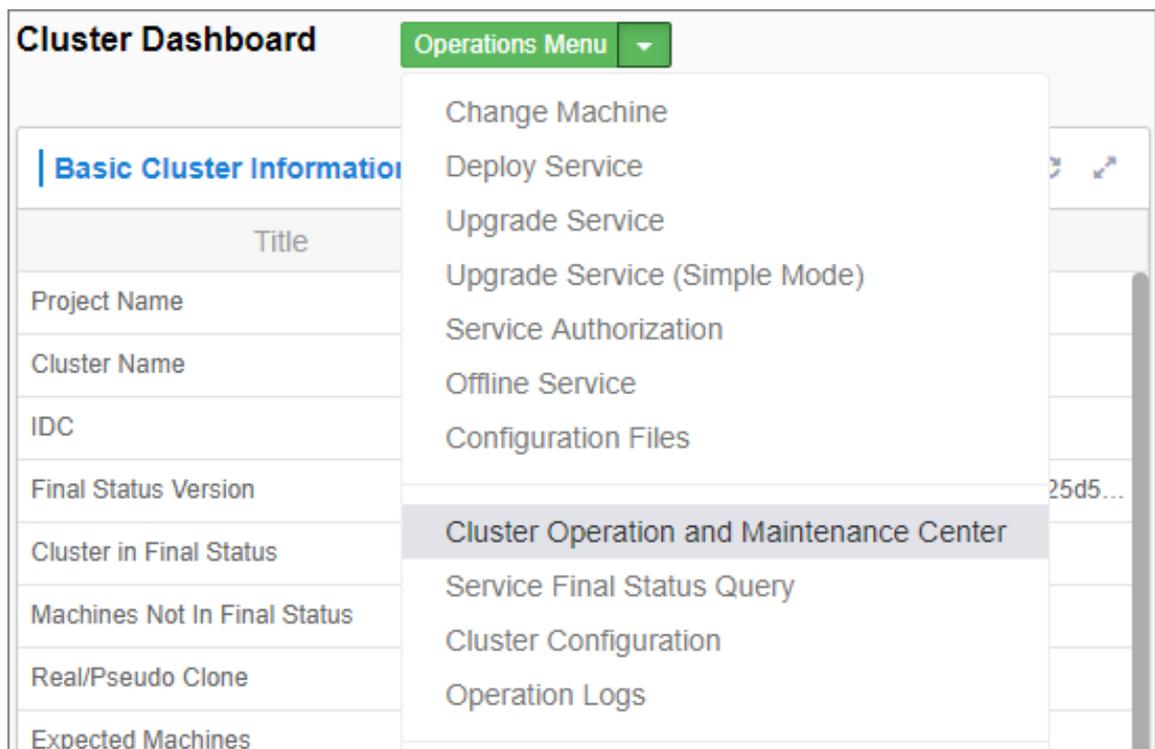
If a green check mark is displayed after dnsServerRole#, then dnsServerRole# is at desired state.

Figure 2-48: View API status

dnsProduct	Final	4 Days 19 Hours	Cluster: 2 / 2	Service: 9 / 9	Role: 12 / 12	Details
drds	Final	4 Days 7 Hours	dnsCluster-A-20...	dnsService	ServiceTest#	
dts	Final	3 Days 23 Hours	standardCluster-...	hids-client	bindServerRole#	
ecs	Final	1 Hour 24 Minutes		os	dnsServerRole#	
edas	Final	4 Days 21 Hours		tianji	dnsServiceDbInit#	
elasticsearch	Final	11 Hours 57 Minutes		tianji-dockerdae...	monitorSrDemo#	
emr	Final	4 Days 21 Hours				
ess	Final	3 Days 22 Hours				

2. Obtain the IP addresses of servers where the API services are deployed.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose Operations > Cluster Operations.
 - c) Click a cluster URL to open the Cluster Dashboard page.
 - d) On the Cluster Dashboard page, choose Operations Menu > Cluster Operation and Maintenance Center.

Figure 2-49: Cluster Operation and Maintenance Center



- e) On the Cluster Operation and Maintenance Center page, view and obtain the IP addresses of servers that are deployed with the API service.

Figure 2-50: View the IP addresses of servers

The screenshot displays the 'Cluster Operation and Maintenance Center' for a cluster named 'dnsCluster-A-20190827-4eb3'. It shows a summary of server status with 'Total Machines' at 2 and 'Abnormal Machines' at 0. Below this, a table lists individual machines with their IP addresses, SR status, and actions.

Machine	Final SR Status	SR Running Status	Action	Action Status	Monitoring Statistics	Actions
vm010012012075 10.12.13.15	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 7	Terminal, Approval Action, Restart Server Role
vm010012016048 10.12.13.48	Normal	GOOD	N/A	N/A	Error: 0 Warning: 0 Good: 7	Terminal, Approval Action, Restart Server Role

3. Log on to the DNS API server. Run the `curl http://localhost/checkpreload.htm` command, and check whether the command output is "success".
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose Operations > Server Operations.
 - c) Click Terminal in the Actions column of a server to log on to the server.
 - d) Run the `curl http://localhost/checkpreload.htm` command on the server where the API service is deployed and check whether the command output is "success".

Figure 2-51: Verify the server

```
[root@docker010036017163 ~]# curl localhost/checkpreload.htm
success
```

2.14.3.2 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.

2.14.4 DNS system

2.14.4.1 Check whether the service role is normal

Procedure

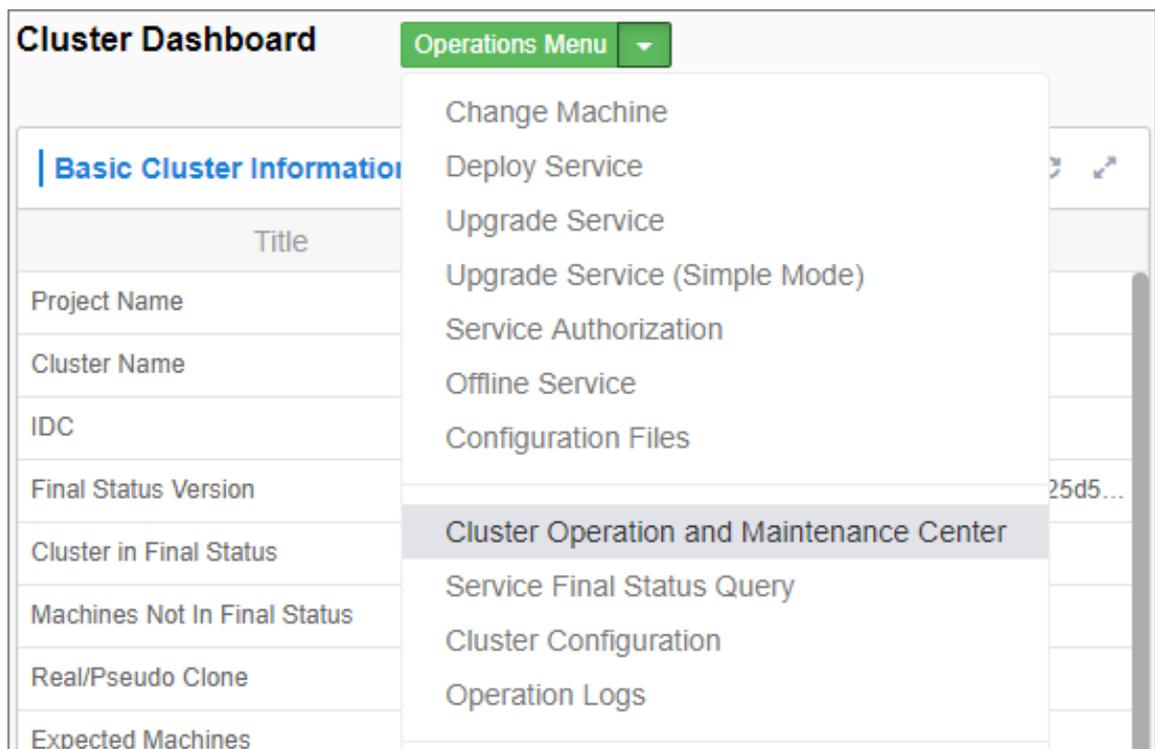
1. You can check whether Apsara Stack DNS is at desired state in the Apsara Infrastructure Management Framework console.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose **Tasks > Deployment Summary** to open the **Deployment Summary** page.
 - c) Click **Deployment Details**.
 - d) On the **Deployment Details** page, find the `dnsProduct` project.
 - e) Find the `bindServerRole#` service role, and click **Details** in the **Deployment Progress** column to check whether the service role is at desired state.

Figure 2-52: Check whether the `bindServerRole#` service role 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	dnsCluster-A-20...	dnsService	ServiceTest#	✓
dtc	Final	3 Days 23 Hours	standardCluster-...	hids-client	bindServerRole#	✓
ecs	Final	1 Hour 24 Minutes		os	dnsServerRole#	✓
edas	Final	4 Days 21 Hours		tianji	dnsServiceDbInit#	✓
elasticsearch	Final	11 Hours 57 Minutes		tianji-dockerdae...	monitorSrDemo#	✓
emr	Final	4 Days 21 Hours				
ess	Final	3 Days 22 Hours				

2. Obtain the IP addresses of servers where the DNS services are deployed.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose Operations > Cluster Operations.
 - c) Click a cluster URL to open the Cluster Dashboard page.
 - d) On the Cluster Dashboard page, choose Operations Menu > Cluster Operation and Maintenance Center.

Figure 2-53: Cluster Operation and Maintenance Center



- e) On the Cluster Operation and Maintenance Center page, view and obtain all server IP addresses for which bindServerRole# provides services.

3. Log on to the DNS server, and check whether the status code returned by `python /bind/hello/check_health.py|echo $?` is 0.
 - a) Log on to the Apsara Infrastructure Management Framework console.
 - b) In the top navigation bar, choose **Operations > Server Operations**.
 - c) Click **Terminal** in the **Actions** column of a server to log on to the server.
 - d) Run the `python /bind/hello/check_health.py|echo $?` command on the server where `bindServerRole#` is deployed and check whether the status code 0 is returned.

Figure 2-54: Verify the server

```
[root@101h08207.cloud.h10.amtest1284 /bind/hello]
#python check_health.py|echo $?
0
```

2.14.4.2 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.

2.14.4.3 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.

2.14.5 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.**

2.14.6 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.**

2.15 API Gateway

2.15.1 API Gateway introduction

This topic describes Apsara Stack API Gateway and the features of its modules.

API Gateway console

The API Gateway console is used to configure and manage your APIs and related policies. With the API management system, you can query, update, edit, and delete APIs. You can also create, associate, disassociate, and delete API management policies. API Gateway also provides a full range of API lifecycle management functions, including creating, testing, publishing, and unpublishing APIs. It improves API management and iteration efficiency. All your data will eventually be used as the API metadata for API Gateway.

API Gateway

API Gateway is a complete API hosting service. It helps you use APIs to provide capabilities, services, and data to your partners. API Gateway is initialized based on the API metadata generated by the API management system, and ultimately acts as the agent to send API requests. API Gateway provides a range of mechanisms to enhance security and reduce risks arising from APIs. These mechanisms include attack prevention, replay prevention, request encryption, identity authentication, permission management, and throttling.

2.15.2 Routine maintenance

2.15.2.1 View operation logs

During O&M, if you need to view logs to troubleshoot errors, you can query logs stored in relevant locations on different systems.

API Gateway OpenAPI logs: The operation log files are stored in the `/alidata/www/logs/java/cloudapi-openapi/` directory. You can query the files as required.

API Gateway logs: The operation log files are stored in the `/alidata/logs/` directory. Each log contains log entries of a day. You can query the files as required.

2.15.2.2 Enable and disable API services

You can log on to API servers as an administrator and run the `sudo /etc/init.d/jetty restart` command to restart API services. To ensure the continuity of services, we recommend that you do not restart the API services on all servers at the same time.

You can run the `sudo sh /home/admin/stop.sh` command to stop API services, and run the `sudo sh /home/admin/start.sh` command to start API services.

2.15.3 API Gateway O&M

2.15.3.1 System O&M

2.15.3.1.1 Check the desired state of API Gateway

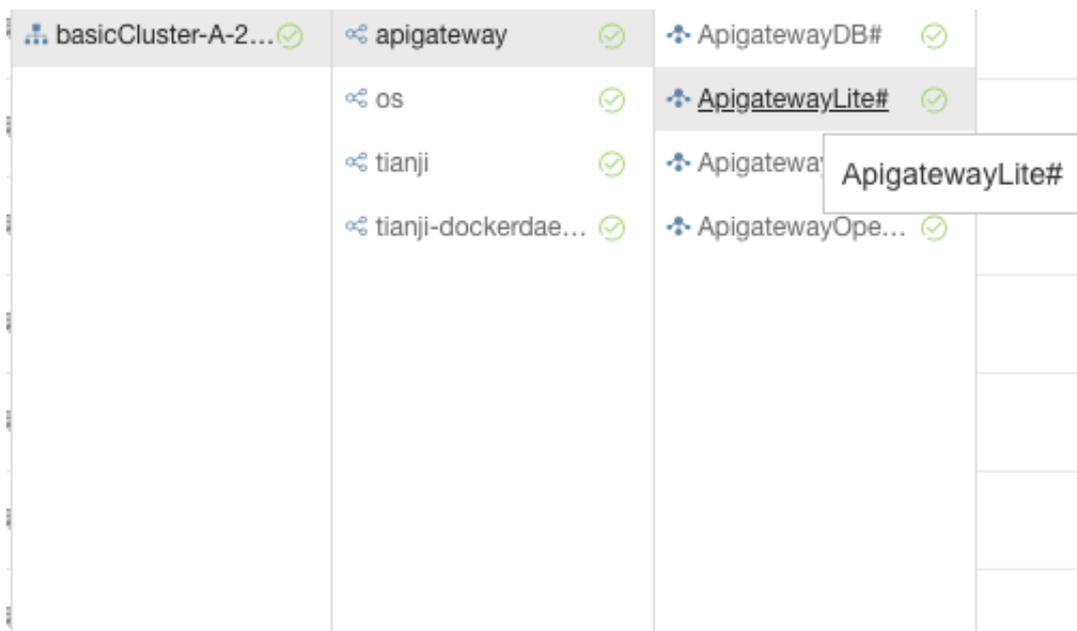
You can use Apsara Infrastructure Management Framework to operate and maintain API Gateway. To log on to the machines in which the API Gateway console

resides, choose **Operations > Server Operations** in the **Apsara Infrastructure Management Framework** console.

Procedure

1. Log on to the **Apsara Infrastructure Management Framework** console.
2. In the top navigation bar, choose **Tasks > Deployment Summary**.
3. On the **Deployment Summary** page that appears, click **Deployment Details**.
4. On the **Deployment Details** page, find the **apigateway** project.
5. Click **Details** in the **Deployment Progress** column corresponding to the **apigateway** project. Check whether the **ApigatewayLite#** server role is in the desired state.

If a green tick appears for the server role item, the server role has reached the desired state.

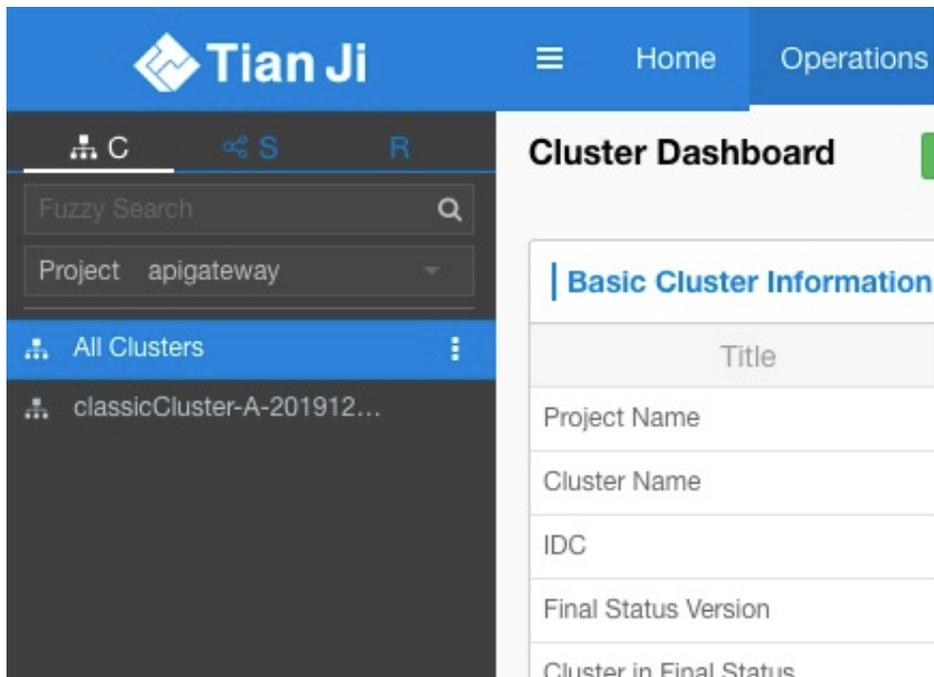


2.15.3.1.2 Check the service status of OpenAPI

Procedure

1. Find machines in the ApigatewayOpenAPI# server role.

- a) Log on to the Apsara Infrastructure Management Framework console.
- b) Click the C tab in the left-side navigation pane.
- c) Select apigateway from the Project drop-down list.



- d) Place the pointer over the  icon next to one of the filtered clusters and choose Dashboard from the shortcut menu.
- e) In the Service Instance List section, click Details in the Actions column corresponding to the apigateway service instance.
- f) In the Server Role List section, you can view the deployment status of each role.

Server Role	Current Status	Expected Machines	Machines In Final...	Machines Going ...	Rolling Task Status	Time Used	Actions
ApigatewayConsole#	In Final Status	2	2	0	no rolling		Details
ApigatewayDB#	In Final Status	1	1	0	no rolling		Details
ApigatewayLite#	In Final Status	3	3	0	no rolling		Details
ApigatewayOpenAPI#	In Final Status	2	2	0	no rolling		Details
ServiceTest#	In Final Status	1	1	0	no rolling		Details

- g) Click Details in the Actions column corresponding to the ApigatewayOpenAPI# role and view machine information of the role in the Machine Information section.

Machine Information									
Mac...	IP	Machi...	Machi...	Server...	Server...	Curren...	Target ...	Error ...	Actions
vm01001...	10.11.106...	good		good PR...		f52a09921...	f52a09921...		Terminal Restart Details Machine System View Machine Operation
vm01001...	10.11.106...	good		good PR...		f52a09921...	f52a09921...		Terminal Restart Details Machine System View Machine Operation

2. Click **Terminal** in the **Actions** column corresponding to a machine to log on to the machine.

3. Run the following command to find the container:

```
docker ps|grep cloudapi-openapi
```

4. Run the following command to find the container IP address:

```
docker inspect [container ID] | grep IPAddress
```

5. Run the following command to check whether OK is returned:

```
curl -i http://localhost:18080/cloudapi-openapi/check_health
```

```
[admin@vm010148065157 /home/admin]
$docker inspect 81e002d83e7b |grep IPAddress
      "SecondaryIPAddresses": null,
      "IPAddress": "",
      "IPAddress": "10.148.65.158",

[admin@vm010148065157 /home/admin]
$curl http://10.148.65.158:18080/cloudapi-openapi/check_health
ok
```

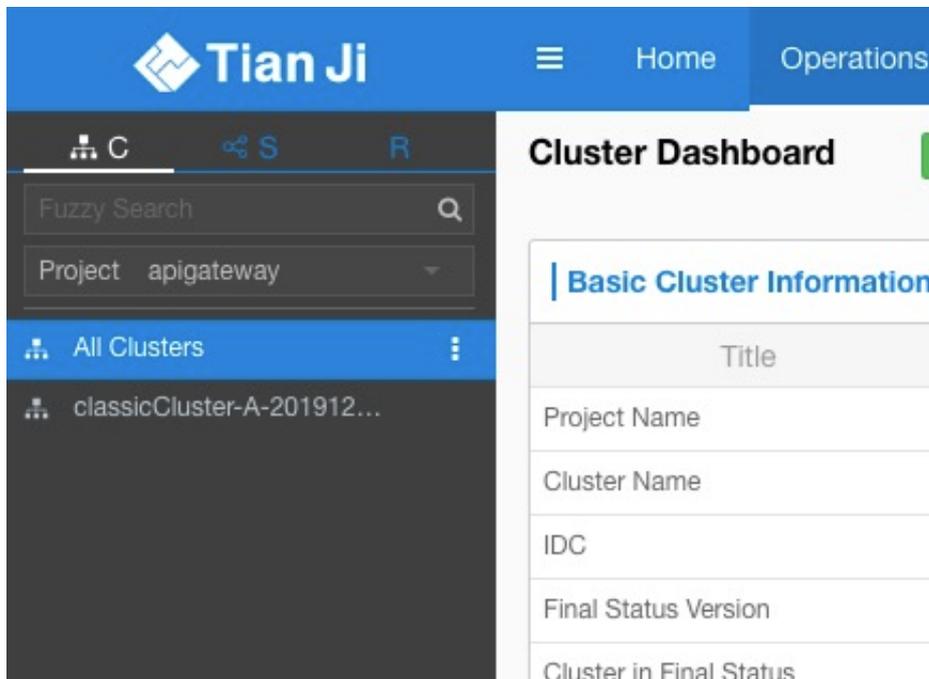
If OK is returned, the service status of the OpenAPI component is normal.

2.15.3.1.3 Check the service status of the API Gateway console

Procedure

1. Find machines in the ApigatewayConsole# server role.

- a) Log on to the Apsara Infrastructure Management Framework console.
- b) Click the C tab in the left-side navigation pane.
- c) Select apigateway from the Project drop-down list.



- d) Place the pointer over the  icon next to one of the filtered clusters and choose Dashboard from the shortcut menu.
- e) In the Service Instance List section, click Details in the Actions column corresponding to the apigateway service instance.
- f) In the Server Role List section, you can view the deployment status of each role.

Server Role	Current Status	Expected Machines	Machines In Final...	Machines Going ...	Rolling Task Status	Time Used	Actions
ApigatewayConsole#	In Final Status	2	2	0	no rolling		Details
ApigatewayDB#	In Final Status	1	1	0	no rolling		Details
ApigatewayLite#	In Final Status	3	3	0	no rolling		Details
ApigatewayOpenAPI#	In Final Status	2	2	0	no rolling		Details
ServiceTest#	In Final Status	1	1	0	no rolling		Details

- g) Click Details in the Actions column corresponding to the ApigatewayConsole# role and view machine information of the role in the Machine Information section.

Machine Information									
Mac...	IP	Machi...	Machi...	Server...	Server...	Curren...	Target ...	Error ...	Actions
vm01001...	10.11.106...	good		good PR...		f52a09921...	f52a09921...		Terminal Restart Details Machine System View Machine Operation
vm01001...	10.11.106...	good		good PR...		f52a09921...	f52a09921...		Terminal Restart Details Machine System View Machine Operation

2. Click **Terminal** in the **Actions** column corresponding to a machine to log on to the machine.

3. Run the following command to find the container:

```
docker ps|grep cloudapi-openapi
```

4. Run the following command to find the container IP address:

```
docker inspect [container ID] | grep IPAddress
```

5. Run the following command to check whether OK is returned:

```
curl -i http://localhost:18080/cag-console-aliyun-com/check_health
```

```
[admin@vm010148065157 /home/admin]
$docker ps|grep console-backend
bc0d1d8295ea        696fa22ae150        "/bin/sh -c /alidata/"   3 days ago         Up 3 days          apigateway.ApigatewayConsole_
.console-backend.1566551509

[admin@vm010148065157 /home/admin]
$docker inspect bc0d1d8295ea|grep IPAddress
      "SecondaryIPAddresses": null,
      "IPAddress": "",
      "IPAddress": "10.148.65.159",

[admin@vm010148065157 /home/admin]
$curl http://10.148.65.159:18080/cag-console-aliyun-com/check_health
ok
[admin@vm010148065157 /home/admin]
```

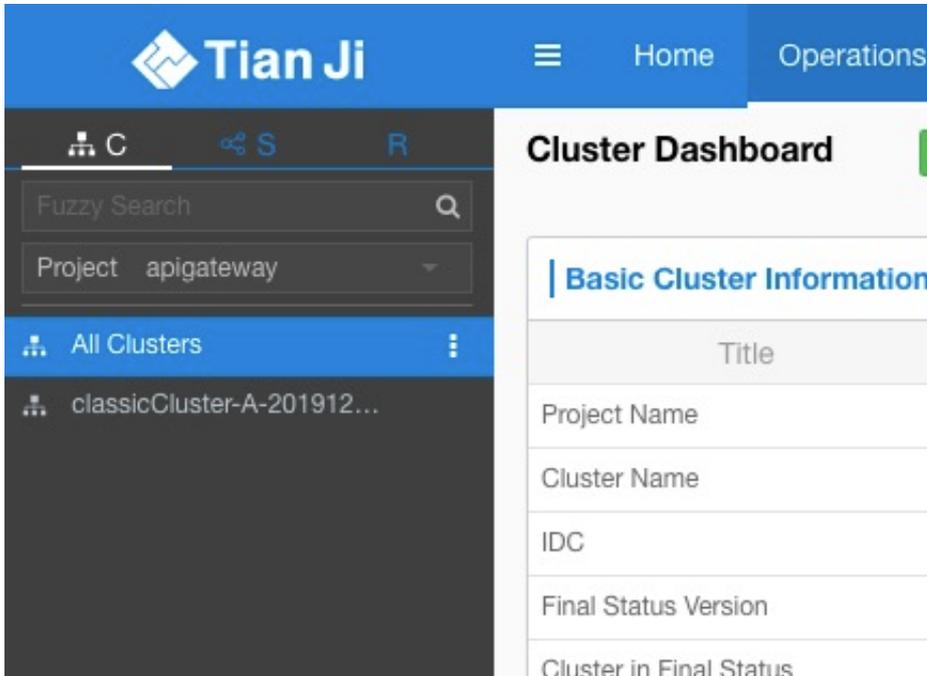
If OK is returned, the service status of the API Gateway console is normal.

2.15.3.1.4 Check the service status of API Gateway

Procedure

1. Find machines in the ApigatewayLite# server role.

- a) Log on to the Apsara Infrastructure Management Framework console.
- b) Click the C tab in the left-side navigation pane.
- c) Select apigateway from the Project drop-down list.



- d) Place the pointer over the  icon next to one of the filtered clusters and choose Dashboard from the shortcut menu.
- e) In the Service Instance List section, click Details in the Actions column corresponding to the apigateway service instance.
- f) In the Server Role List section, you can view the deployment status of each role.

Server Role	Current Status	Expected Machines	Machines In Final...	Machines Going ...	Rolling Task Status	Time Used	Actions
ApigatewayConsole#	In Final Status	2	2	0	no rolling		Details
ApigatewayDB#	In Final Status	1	1	0	no rolling		Details
ApigatewayLite#	In Final Status	3	3	0	no rolling		Details
ApigatewayOpenAPI#	In Final Status	2	2	0	no rolling		Details
ServiceTest#	In Final Status	1	1	0	no rolling		Details

- g) Click Details in the Actions column corresponding to the ApigatewayLite# role and view machine information of the role in the Machine Information section.

Mac...	IP	Machi...	Machi...	Server...	Server...	Curren...	Target ...	Error ...	Actions
vm01001...	10.11.106...	good		good PR...		f52a09921...	f52a09921...		Terminal Restart Details Machine System View Machine Operation
vm01001...	10.11.106...	good		good PR...		f52a09921...	f52a09921...		Terminal Restart Details Machine System View Machine Operation

2. Click **Terminal** in the **Actions** column corresponding to a machine to log on to the machine.
3. Run the following command to check whether the I'm fine, thank you, and you? message is returned:

```
curl -i http://localhost/status -H Host:status.taobao.com
```

2.15.3.1.5 View results of automated test cases

Procedure

1. Log on to the Apsara Infrastructure Management Framework console.
2. Click the **C** tab in the left-side navigation pane.
3. Select **apigateway** from the **Project** drop-down list.
4. Place the pointer over the  icon next to one of the filtered clusters and choose **Dashboard** from the shortcut menu.
5. In the **Service Instance List** section, click **Details** in the **Actions** column corresponding to the apigateway service instance.
6. In the **Service Monitoring Information** section, click **Details** in the **Actions** column to view the automated test case report.

Service Monitoring Information				
Monitored Item	Level	Description	Updated At	Actions
test_report	info	{"name":"cloudapi-...	01/12/20, 11:07:13	Details

2.15.3.2 Troubleshooting

Context



Note:

- `/alidata/logs/system.log`: API Gateway logs.
- `/usr/share/jetty/logs/stderrout.log`: API Gateway console and OpenAPI logs.

Procedure

Start the application and check whether any errors have occurred. Check whether the system is operating normally.

- **If the system is operating but does not function properly, check the logs to troubleshoot errors.**
- **If the system quits shortly after being started up, check the logs to troubleshoot errors.**

2.15.4 Log analysis

You can perform log analysis based on the ID of an individual API request.

After you send a request, you will receive a response that contains the request ID from API Gateway.

You can use the request ID to perform the following operations:

- **All API Gateway logs are uploaded to Log Service, where you can view the request ID.**
- **You can use the request ID to query the response to or error message for the current request in the API system logs.**

3 Operations of big data products

3.1 Apsara Bigdata Manager (ABM) platform

3.1.1 What is Apsara Bigdata Manager?

Apsara Bigdata Manager (ABM) is an O&M platform for big data products, including MaxCompute, DataWorks, StreamCompute, Quick BI, Graph Analytics, Elasticsearch, Dataphin, DataHub, and Machine Learning Platform for AI.

ABM supports O&M on the business, services, clusters, and hosts of these big data products. Besides, you can upgrade big data products, customize alert configurations, and view the O&M history in ABM.

By using ABM, on-site Apsara Stack engineers can easily manage big data products, such as viewing resource usage, checking alerts and fix methods, and modifying configurations.

3.1.2 Common operations

The data tables and legends in the Apsara Bigdata Manager (ABM) console facilitate operations. This topic uses MaxCompute and DataHub as examples to describe the common operations.

Search for a project quickly

You can quickly search for a project based on the project name.

- 1. On the MaxCompute page, click O&M in the upper-right corner, and then click Business. The Project List page under Projects appears.**
- 2. In the Quick Search field, enter the project name. Auto-suggestion is supported. Select the target project from the drop-down list, or select the project by using the up and down arrow keys, and then press Enter.**



Note:

When a project is matched, the region of the project appears before the project name.

Quick Search: admin

Filter: cn-... admin_tas...

Project	Cluster	Quota Group	Physical Storage	Logical Storage	File Count	Jobs	Owner	Created At
aaaodps	HYBRIDODPSCLUSTER-A-2	QuotaGroup95eb6831556	14.32 M	4.77 M	2971		ALYUN...	2019-04-30 09:23:17
admin_task_project	HYBRIDODPSCLUSTER-A-2	odps_quota	3.58 K	1.19 K	1		ALYUN...	2019-03-05 00:03:47
ads	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN...	2019-03-05 00:10:41
adsnr	HYBRIDODPSCLUSTER-A-2	BCCDTCENTERAPITESTCRE	25.24 M	8.41 M	2157	8	ALYUN...	2019-03-05 00:10:41
alqo_market	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN...	2019-06-21 00:06:14

Example:

Quick Search: cn-... admin_task_

Filter: cn-... admin_task_

Project	Cluster	Quota Group	Physical Storage	Logical Storage	File Count	Jobs	Owner	Created At	Description	Actions
admin_task_project	HYBRIDODPSCLUSTER-A-2	odps_quota	3.58 K	1.19 K	1		ALYUN...	2019-03-05 00:03:47		Modify Copy-Resource

1 to 1 of 1

Filter projects

You can set filter conditions for multiple columns at the same time to quickly filter the items you want.

1. On the MaxCompute page, click O&M in the upper-right corner, and then click Business. The Project List page under Projects appears.
2. On the Project List page, click Filter in the upper-left corner of the list. A field for setting filter conditions appears for each column.

3. Click the icon next to each field for setting filter conditions and select the filtering method. The default method is Contains.

Project	Cluster	Quota Group	Physical Storage	Logical Storage	File Count	Jobs	Owner
aaodps		taGroup95eb6831556	14.32 M	4.77 M	2971		ALYUN\$
admin_task_project		s_quota	3.58 K	1.19 K	1		ALYUN\$
ads		ps_quota	0	0	0		ALYUN\$
adsmr		CDTCENTERAPITESTCR	25.24 M	8.41 M	2157	8	ALYUN\$
algo_market	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$
algo_public	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$

Optional filtering methods include:

- Equals
- Not equal
- Starts with
- Ends with
- Contains
- Not contains

4. After you select the filtering method, enter the filter condition. The items that meet the filter condition are automatically filtered.

Project	Cluster	Quota Group	Physical Storage	Logical Storage	File Count	Jobs	Owner	Created At	Description	Actions
ad		s_quota	3.58 K	1.19 K	1		ALYUN\$	2019-03-05 00:03:47		Modify Copy-Resource
ads		ps_quota	0	0	0		ALYUN\$	2019-03-05 00:10:41		Modify Copy-Resource
adsmr	HYBRIDODPSCLUSTER-A-2	BCCDTCENTERAPITESTCR	25.24 M	8.41 M	2157	8	ALYUN\$	2019-03-05 00:10:41		Modify Copy-Resource
bigdatademo	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$	2019-04-24 18:52:10		Modify Copy-Resource

5. If the filtering result does not meet the requirements, you can continue setting filter conditions for other columns.

Project	Cluster	Quota Group	Physical Storage	Logical Storage	File Count	Jobs	Owner	Created At	Description	Actions
ad		odps	3.58 K	1.19 K	1		ALYUN\$	2019-03-05 00:03:47		Modify Copy-Resource
ads	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$	2019-03-05 00:10:41		Modify Copy-Resource
bigdatademo	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$	2019-04-24 18:52:10		Modify Copy-Resource

After you set the filter conditions for the projects, the Filter button is highlighted. If you need to cancel filtering, click the highlighted Filter button.

Customize a column

You can customize columns in the list. For example, you can set the column position or column width, and determine whether to display a column. You can also set filter conditions for columns.

On the Project List page, you can drag a column to change its position.

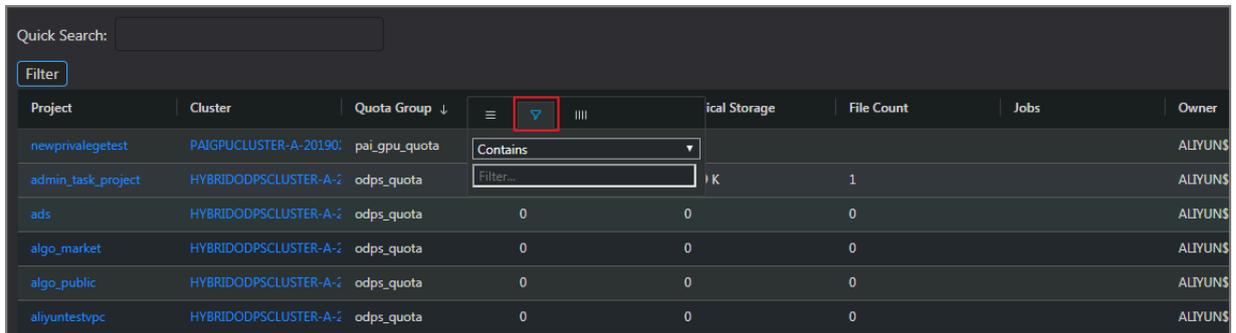
Project	Cluster	Quota Group	Physical Storage	Storage	File Count	Jobs	Owner
ads	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$
algo_market	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$
algo_public	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$
aliyuntestvpc	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0		ALYUN\$
base_1	HYBRIDODPSCLUSTER-A-2	QuotaGroup8102aa61561f	0	0	0		ALYUN\$
base_test01_dev	HYBRIDODPSCLUSTER-A-2	BCCDTCENTERAPITESTCRE	0	0	0		ALYUN\$

You can click  in a column heading to customize the column.

Project	Cluster	Quota Group ↓	Physical Storage	Storage	File Count
newprivalegetest	PAIGPUCLUSTER-A-20190	pai_gpu_quota			
admin_task_project	HYBRIDODPSCLUSTER-A-2	odps_quota			1
ads	HYBRIDODPSCLUSTER-A-2	odps_quota			0
algo_market	HYBRIDODPSCLUSTER-A-2	odps_quota			0
algo_public	HYBRIDODPSCLUSTER-A-2	odps_quota			0
aliyuntestvpc	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0
base_meta	HYBRIDODPSCLUSTER-A-2	odps_quota	371.28 G	123.76 G	33230
bigdatademo	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0
cosmo_pully	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0
dataphin_meta	HYBRIDODPSCLUSTER-A-2	odps_quota	89.62 M	29.87 M	978

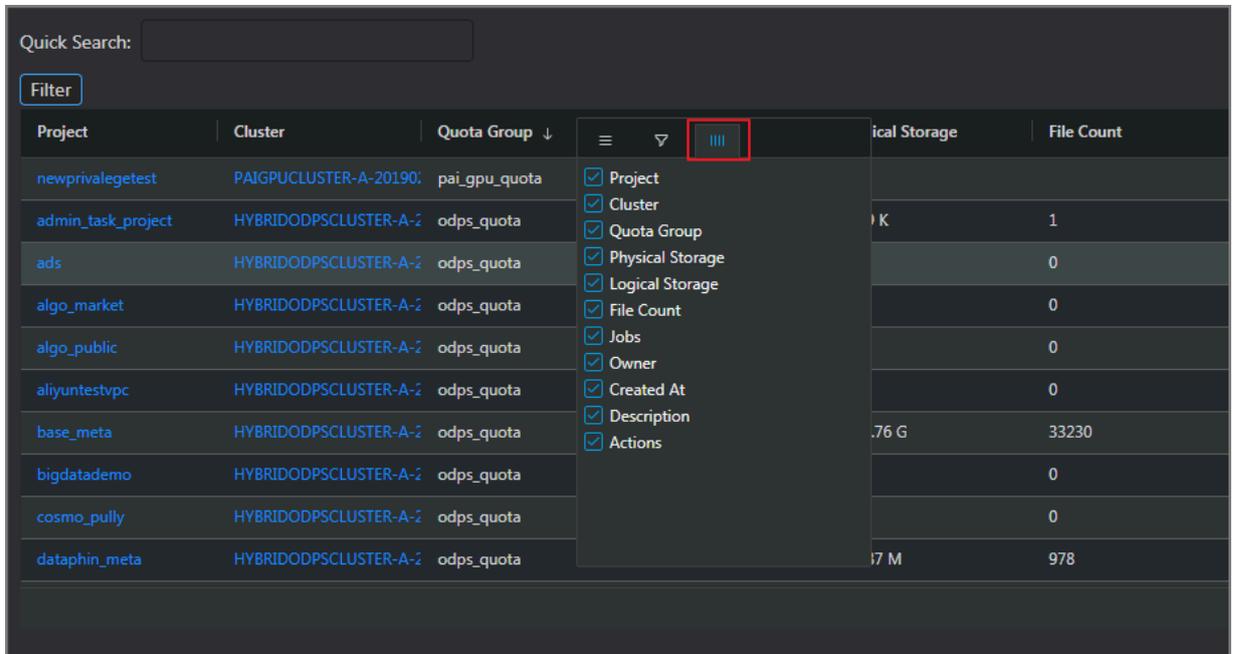
- **Pin Column:** allows you to fix a column to the rightmost or leftmost of the list. Unless being pinned, a column appears at the default position.
- **Autosize This Column:** allows you to adjust the width of a column automatically.
- **Autosize All Columns:** allows you to adjust the width of all columns automatically.
- **Reset Columns:** allows you to reset a column to its initial status.
- **Tool Panel:**

Click  in a column heading and set a filter condition to filter projects based on the column.



Project	Cluster	Quota Group ↓		ical Storage	File Count	Jobs	Owner
newprivalegetest	PAIGPUCLUSTER-A-20190	pai_gpu_quota	Contains				ALIYUN\$
admin_task_project	HYBRIDODPSCLUSTER-A-2	odps_quota	Filter...	1 K	1		ALIYUN\$
ads	HYBRIDODPSCLUSTER-A-2	odps_quota		0	0	0	ALIYUN\$
algo_market	HYBRIDODPSCLUSTER-A-2	odps_quota		0	0	0	ALIYUN\$
algo_public	HYBRIDODPSCLUSTER-A-2	odps_quota		0	0	0	ALIYUN\$
aliyuntestvpc	HYBRIDODPSCLUSTER-A-2	odps_quota		0	0	0	ALIYUN\$

Click  to set the columns to be displayed.



Project	Cluster	Quota Group ↓		ical Storage	File Count
newprivalegetest	PAIGPUCLUSTER-A-20190	pai_gpu_quota	<input checked="" type="checkbox"/> Project <input checked="" type="checkbox"/> Cluster <input checked="" type="checkbox"/> Quota Group <input checked="" type="checkbox"/> Physical Storage <input checked="" type="checkbox"/> Logical Storage <input checked="" type="checkbox"/> File Count <input checked="" type="checkbox"/> Jobs <input checked="" type="checkbox"/> Owner <input checked="" type="checkbox"/> Created At <input checked="" type="checkbox"/> Description <input checked="" type="checkbox"/> Actions		
admin_task_project	HYBRIDODPSCLUSTER-A-2	odps_quota		1 K	1
ads	HYBRIDODPSCLUSTER-A-2	odps_quota			0
algo_market	HYBRIDODPSCLUSTER-A-2	odps_quota			0
algo_public	HYBRIDODPSCLUSTER-A-2	odps_quota			0
aliyuntestvpc	HYBRIDODPSCLUSTER-A-2	odps_quota			0
base_meta	HYBRIDODPSCLUSTER-A-2	odps_quota		.76 G	33230
bigdatademo	HYBRIDODPSCLUSTER-A-2	odps_quota			0
cosmo_pully	HYBRIDODPSCLUSTER-A-2	odps_quota			0
dataphin_meta	HYBRIDODPSCLUSTER-A-2	odps_quota		17 M	978

If you select the check box of a column name, the column appears. Otherwise, the column is hidden.

Show the tool panel

After the tool panel appears, it is attached to the right of the list so that you can quickly set the columns to be displayed.

On the Project List page, click  in a column heading and then select Tool Panel.

The tool panel is then attached to the right of the list.

Quick Search:

Filter

Project	Cluster	Quota Group ↓		Physical Storage	File Count
newprivalegetest	PAIGPUCLUSTER-A-20190	pai_gpu_quota	☰		
admin_task_project	HYBRIDODPSCLUSTER-A-2	odps_quota	☰	0 K	1
ads	HYBRIDODPSCLUSTER-A-2	odps_quota	☰		0
algo_market	HYBRIDODPSCLUSTER-A-2	odps_quota	☰		0
algo_public	HYBRIDODPSCLUSTER-A-2	odps_quota	☰		0
aliyuntestvpc	HYBRIDODPSCLUSTER-A-2	odps_quota	☰	0	0
base_meta	HYBRIDODPSCLUSTER-A-2	odps_quota	☰	371.28 G	123.76 G
bigdatademo	HYBRIDODPSCLUSTER-A-2	odps_quota	☰	0	0
cosmo_pully	HYBRIDODPSCLUSTER-A-2	odps_quota	☰	0	0
dataphin_meta	HYBRIDODPSCLUSTER-A-2	odps_quota	☰	89.62 M	29.87 M

Refresh

File Count	Jobs	Owner	Created At	Description
		ALIYUN\$	2019-03-29 18:25:01	
1		ALIYUN\$	2019-03-05 00:03:47	
0		ALIYUN\$	2019-03-05 00:10:41	
0		ALIYUN\$	2019-06-21 00:06:14	
0		ALIYUN\$	2019-03-05 00:10:40	
0		ALIYUN\$	2019-03-26 14:52:12	
33230		ALIYUN\$	2019-03-05 00:10:40	
0		ALIYUN\$	2019-04-24 18:52:10	
0		ALIYUN\$	2019-03-06 18:19:24	
978		ALIYUN\$	2019-03-05 00:10:40	

- Project
- Cluster
- Quota Group
- Physical Storage
- Logical Storage
- File Count
- Jobs
- Owner
- Created At
- Description
- Actions
- Row Groups
Drag here to set row groups
- Values
Drag here to aggregate

1 to 10 of 144 < 1 2 3 4 5 ... 15 >

Sort projects based on a column

The projects can be sorted based on a column in ascending or descending order.

On the Project List page, click a column heading in the list. When you click the column heading for the first time, the projects are sorted based on the column in ascending order. When you click the column heading for the second time, the projects are sorted in descending order. When you click the column heading for the third time, the default sorting is restored.

Quick Search:

Filter

Project ↑	Cluster	Quota Group	Physical Storage	Logical Storage	File Count
aaaodps	HYBRIDODPSCLUSTER-A-2	QuotaGroup95eb6831556f	14.32 M	4.77 M	2971
admin_task_project	HYBRIDODPSCLUSTER-A-2	odps_quota	3.58 K	1.19 K	1
ads	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0
adsmr	HYBRIDODPSCLUSTER-A-2	BCCDTCENTERAPITESTCRE	25.24 M	8.41 M	2157
algo_market	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0
algo_public	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0
aliyuntestvpc	HYBRIDODPSCLUSTER-A-2	odps_quota	0	0	0
base_1	HYBRIDODPSCLUSTER-A-2	QuotaGroup8102aa61561f	0	0	0
base_meta	HYBRIDODPSCLUSTER-A-2	odps_quota	371.28 G	123.76 G	33230
base_test	HYBRIDODPSCLUSTER-A-2	QuotaGroup5f77f1c15532a	3.68 M	1.22 M	24

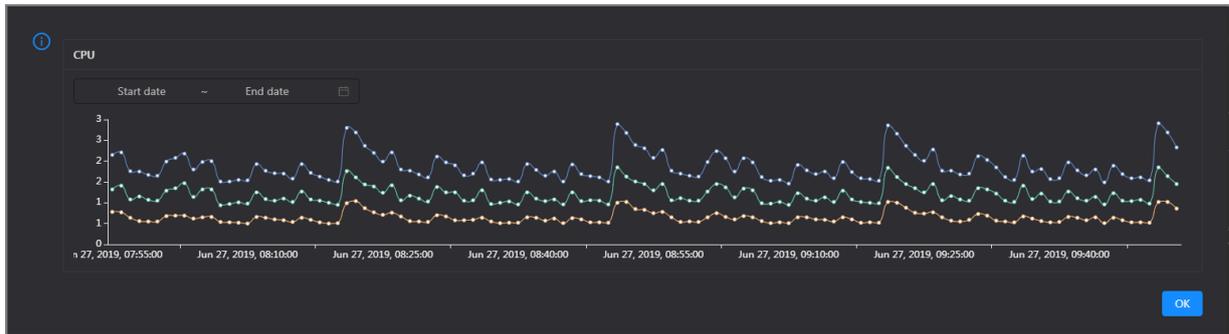
Trend chart 1

On the MaxCompute page, click O&M in the upper-right corner, and then click Clusters. On the Clusters page, you can view relevant metrics, such as CPU and memory, of the selected cluster.



Take CPU as an example. The trend chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) in the specified cluster over time in different colors. You can specify a time period to view the CPU usage.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

Trend chart 2

1. On the DataHub page, click O&M in the upper-right corner, and then click Business.
2. In the right pane of the Business page that appears, select a cluster, and then select a project or a topic in a project. Information about the read and write metrics for the selected project or topic appears in a chart.

The chart displays the trend lines of the read and write metrics for a project or topic over time in different colors. You can click the name of a metric under the chart to determine whether to display the corresponding trend line in the chart. A highlighted metric name indicates that the corresponding trend line is displayed, whereas a dimmed metric name indicates that the corresponding trend line is not displayed.

3.1.3 Quick start

3.1.3.1 Log on to the ABM console

This topic describes how to log on to the Apsara Bigdata Manager (ABM) console.

Context

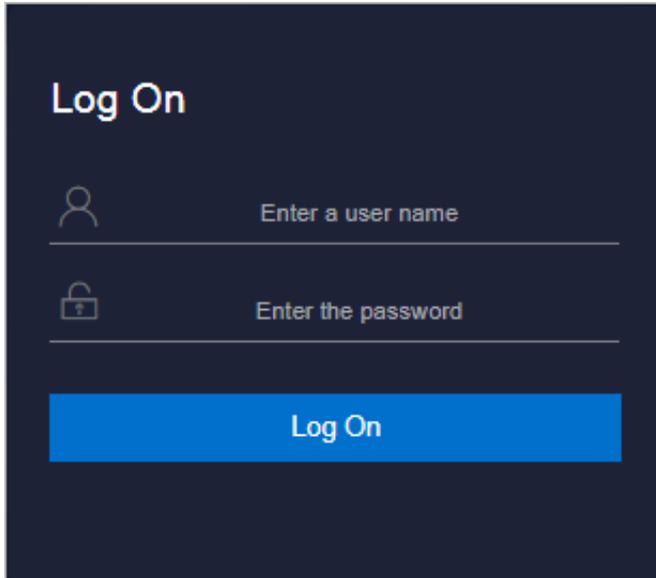
- ASO access address in the format of `http://region-id.aso.intranet-domain-id.com`.
- Google Chrome browser (recommended).

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.

Figure 3-1: Log on to ASO

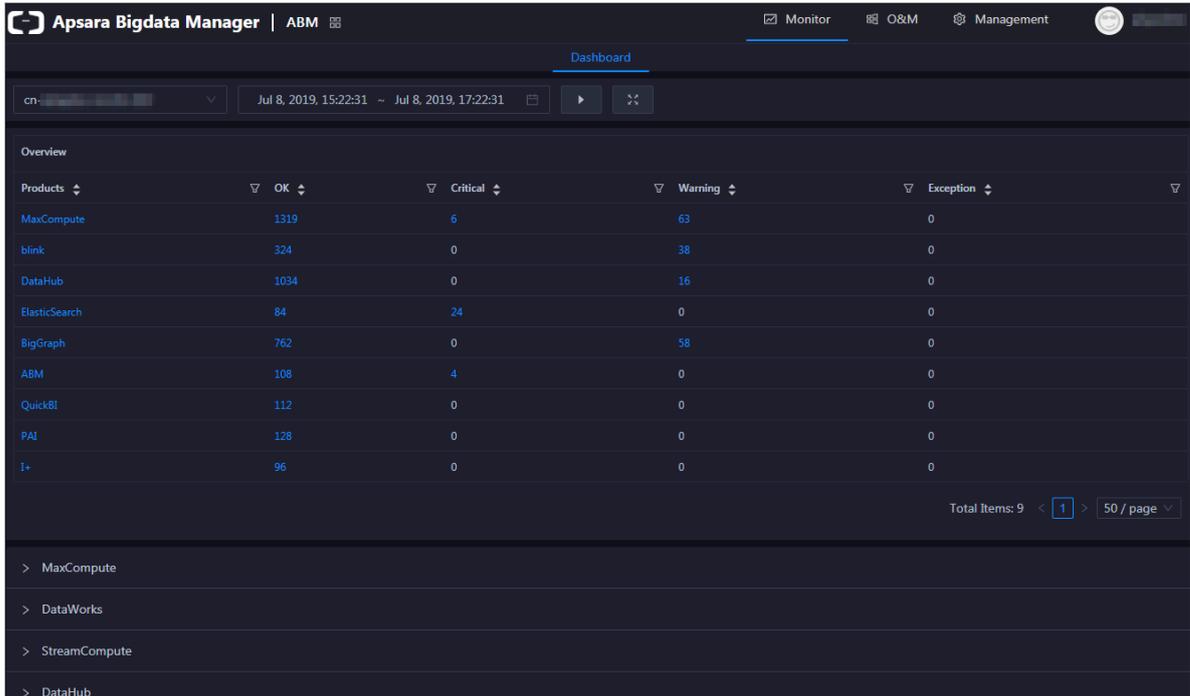


Note:

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

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

5. In the left-side navigation pane, choose **Products > Apsara Bigdata Manager** to log on to the ABM console.



3.1.3.2 Set the background color

You can set the background color of the Apsara Bigdata Manager (ABM) console to white or black based on your preferences. The default color is black.

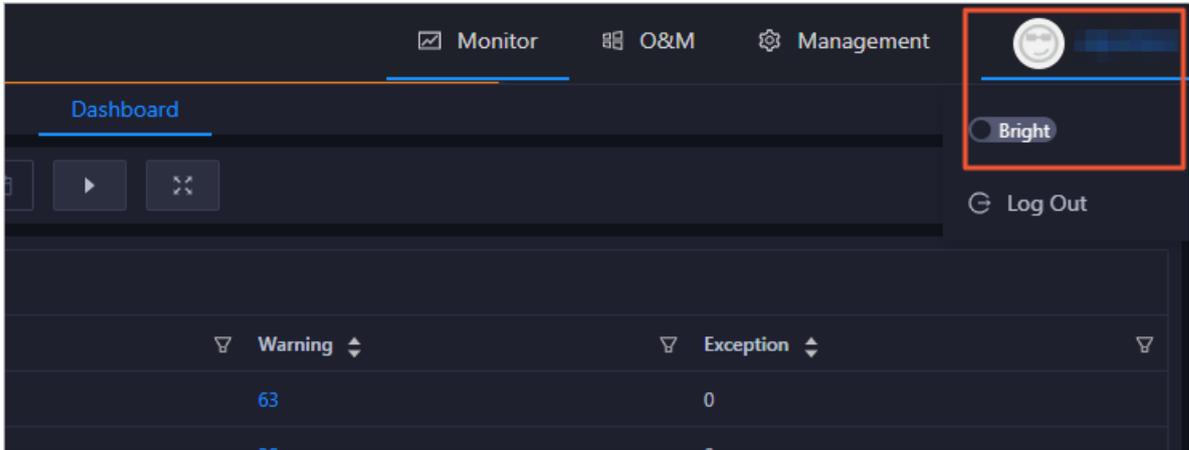
Prerequisites

You have obtained an ABM account and the corresponding password.

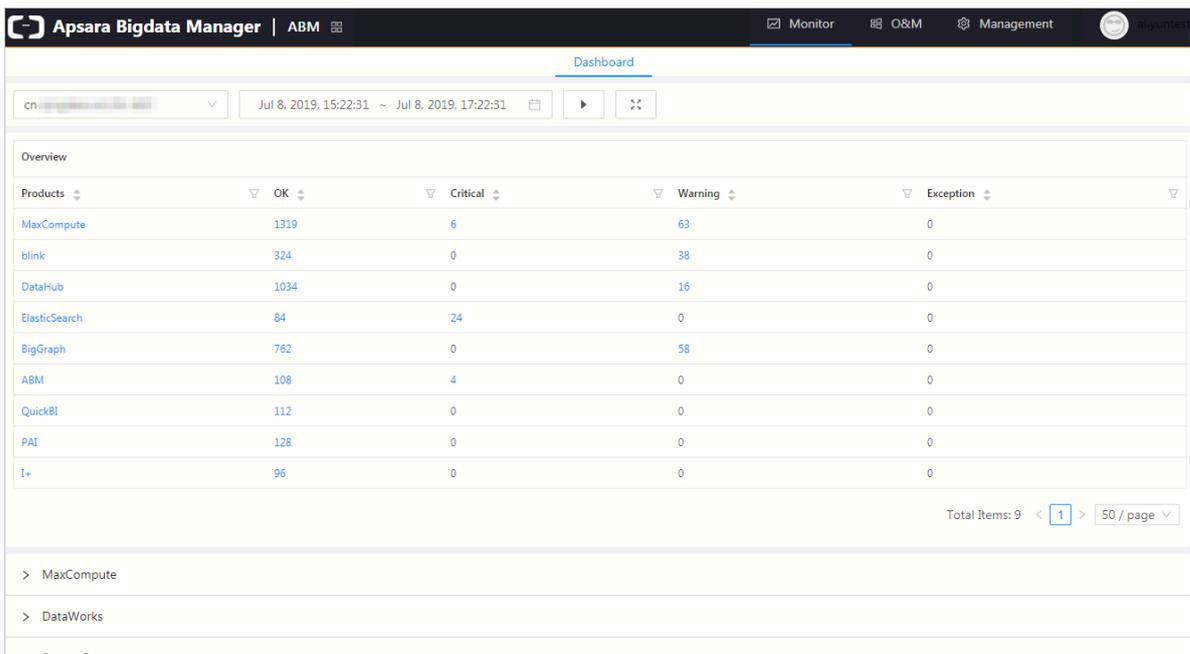
Procedure

1. *Log on to the ABM console.*

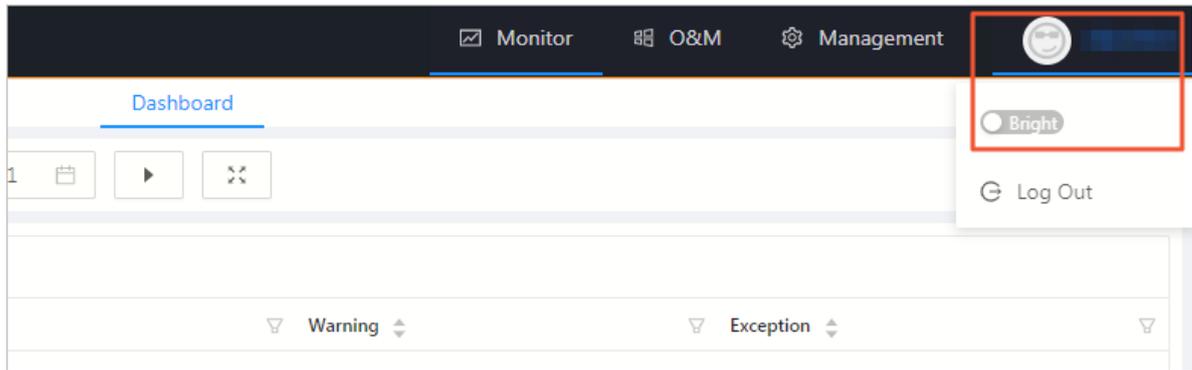
2. Click the avatar in the upper-right corner, and then turn on the Bright to set the background color to white.



The following figure shows the ABM console when the background color is set to white.



To set the background color to black, click the avatar in the upper-right corner, and then turn off the Bright switch.



3.1.3.3 View the dashboard

The Apsara Bigdata Manager (ABM) dashboard displays key operation metrics for the four core services, namely MaxCompute, DataWorks, Realtime Compute, and DataHub. It also provides information about alerts for all big data services, so that you can understand the overall running status of the big data services.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on the corresponding service.

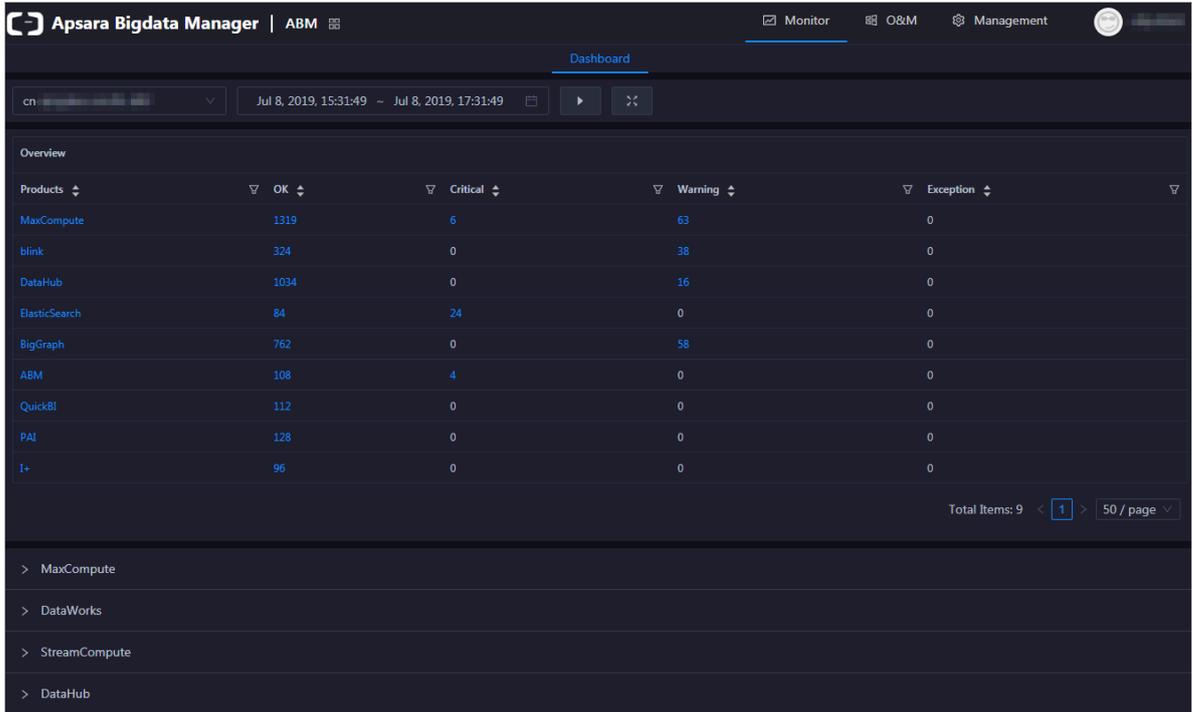
Background

The dashboard is a feature of ABM. As the homepage of ABM, the dashboard allows you to view the overall running information about all big data services.

Procedure

1. *Log on to the ABM console.*

Log on to the ABM console. The Dashboard page appears. To return to the Dashboard page from any other page, click  in the upper-left corner and then click ABM.



The screenshot displays the Apsara Bigdata Manager ABM Dashboard. The top navigation bar includes 'Apsara Bigdata Manager | ABM' and tabs for 'Monitor', 'O&M', and 'Management'. The main content area is titled 'Dashboard' and shows a time range filter for 'Jul 8, 2019, 15:31:49 ~ Jul 8, 2019, 17:31:49'. Below this is an 'Overview' section with a table of product status metrics.

Products	OK	Critical	Warning	Exception
MaxCompute	1319	6	63	0
blink	324	0	38	0
DataHub	1034	0	16	0
ElasticSearch	84	24	0	0
BigGraph	762	0	58	0
ABM	108	4	0	0
QuickBI	112	0	0	0
PAI	128	0	0	0
I+	96	0	0	0

At the bottom of the table, it indicates 'Total Items: 9' and '1 / 50 / page'. Below the table is a sidebar with expandable sections for 'MaxCompute', 'DataWorks', 'StreamCompute', and 'DataHub'.

2. View and clear service alerts.

In the overview section, you can view the number of alerts for each of all big data services. You need to pay special attention to the Critical and Warning alerts. These alerts must be cleared in a timely manner.

- a. On the Dashboard page, click the number of Critical or Warning alerts of a service such as MaxCompute in the overview section. The Health Status page under Clusters of the service appears.

Checker	Source	Critical	Warning	Exception	Actions
eodps_check_nuwa	tcheck	1	0	0	Details
eodps_check_aas	tcheck	1	0	0	Details
bcc_check_ntp	tcheck	0	10	0	Details
eodps_check_schedulerpoolsize	tcheck	0	1	0	Details
bcc_tsar_tcp_checker	tcheck	0	0	0	Details
bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
bcc_host_live_check	tcheck	0	0	0	Details
bcc_process_thread_count_checker	tcheck	0	0	0	Details
bcc_check_load_high	tcheck	0	0	0	Details
bcc_network_tcp_connections_checker	tcheck	0	0	0	Details

On the Health Status page, you can view all the checkers of the service.

- b. Click Details in the Actions column of a checker with alerts. In the Details dialog box that appears, view the details of the checker and the scheme to clear the alerts. Follow the steps in the scheme to clear the alerts.

Name: eodps_check_nuwa **Source:** tcheck
Alias: Nuwa Check **Application:** eodps
Type: system **Scheduling:** Enable
Data Collection: Enable
Default Execution Interval: 0 0/30 * * * ?

Description:

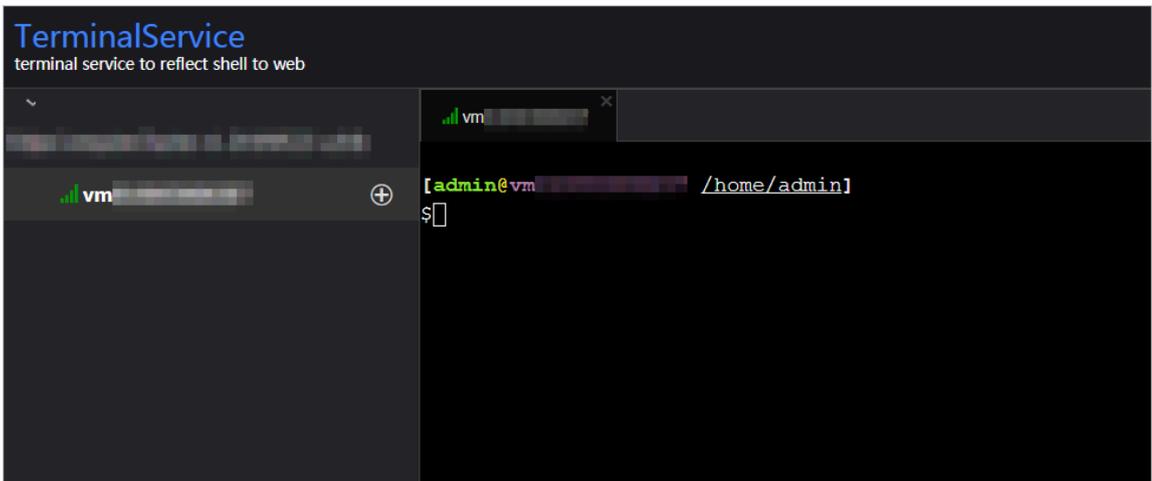
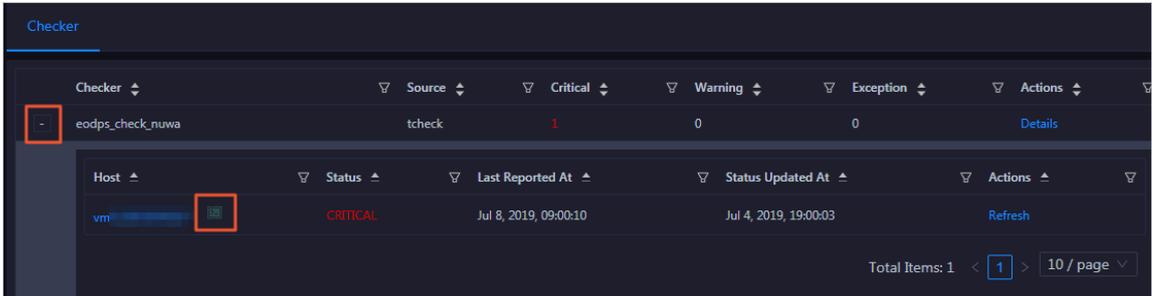
Fix

1. Check the status of Nuwa services: `echo svr[nc localhost 10240]grep Model[grep follower;echo svr[nc localhost 10240]grep Model[grep leader;`
2. Submit a ticket to report the result.

> Show More

- c. Log on to the hosts with alerts if necessary for related operations.

Click + to expand a checker with alerts, and then click the Log On icon next to the name of a host with alerts. On the TerminalService page that appears, click the hostname on the left to log on to the host.

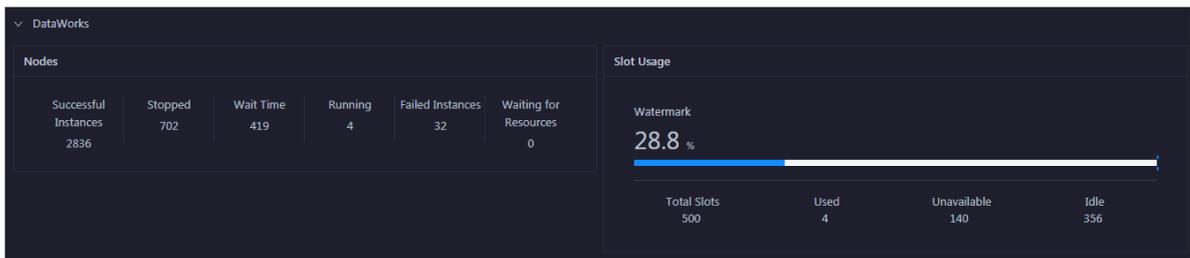


3. On the Dashboard page, click MaxCompute in the overview section to view relevant metrics.



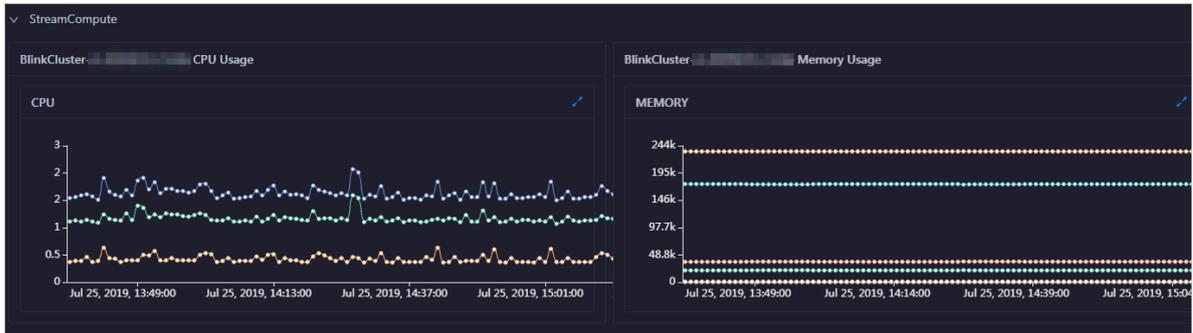
In the MaxCompute section, you can view the allocation of CPU and memory, CPU and memory usage trends, job running status, and storage usage of the MaxCompute cluster.

4. On the Dashboard page, click DataWorks in the overview section to view relevant metrics.



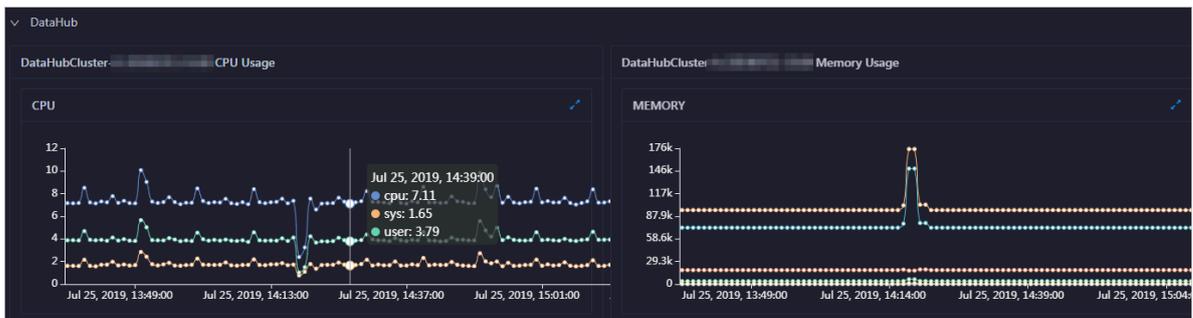
In the DataWorks section, you can view the node scheduling and slot usage of the DataWorks cluster.

5. On the Dashboard page, click StreamCompute in the overview section to view relevant metrics.



In the StreamCompute section, you can view the CPU and memory usage trend charts of the Realtime Compute cluster.

6. On the Dashboard page, click DataHub in the overview section to view relevant metrics.



In the DataHub section, you can view the CPU and memory usage trend charts of the DataHub cluster.

3.1.3.4 View the cluster running status

Apsara Bigdata Manager (ABM) provides you with several operation metrics of clusters, such as CPU usage, memory usage, load, storage, and health check result. This helps you understand the running status of clusters at any time. Based on relevant metrics, you can evaluate whether the selected cluster has operation risks.

Prerequisites

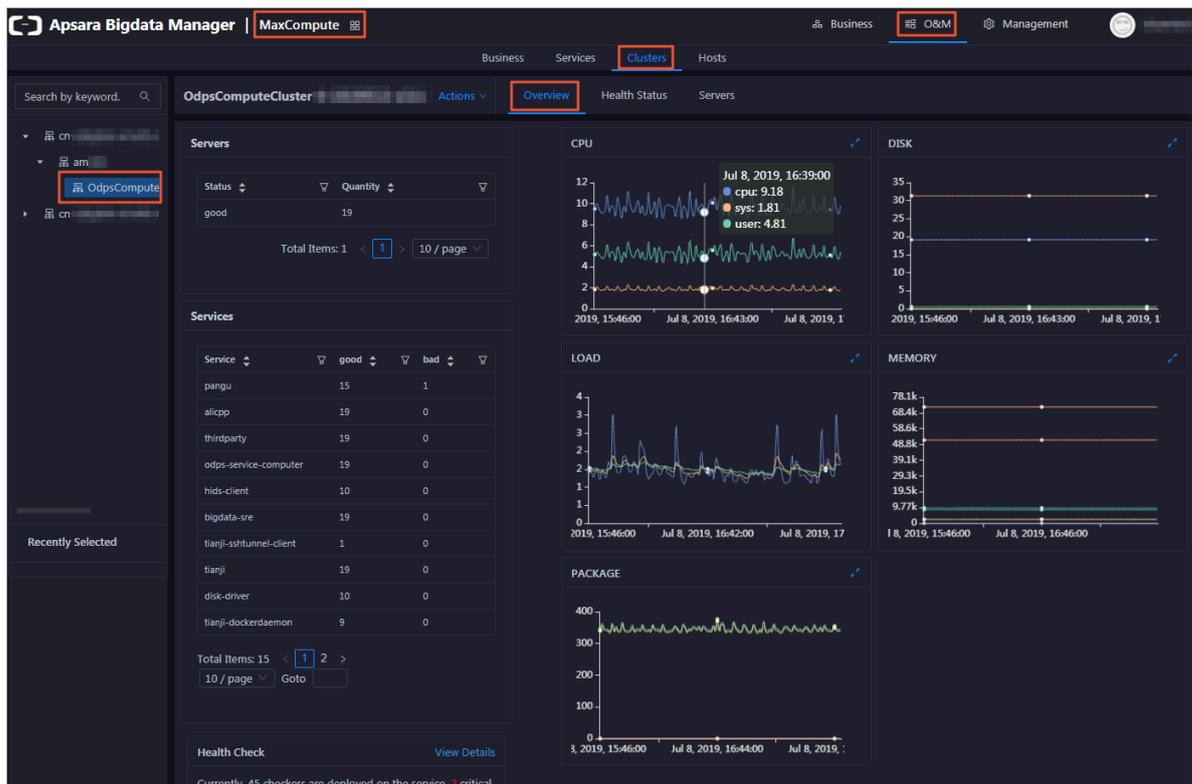
Your ABM account must have the required permissions to perform O&M operations on the corresponding service.

Context

This topic uses MaxCompute as an example to describe how to view the running status of a cluster.

Procedure

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click Clusters.
4. On the Clusters page, select a cluster in the left-side navigation pane. The Overview page for the cluster appears.



On the Overview page, you can view the host status, service status, health check result, and health check history of the selected cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

What's next

You can evaluate the operation risks of a cluster based on the metrics such as the service status, CPU usage, disk usage, memory usage, and load.

If the cluster has any Critical, Warning, or Exception alerts, you need to check and clear them in a timely manner. You need to pay special attention to the Critical and Warning alerts. For more information, see [View and clear cluster alerts.](#)

3.1.3.5 View and clear cluster alerts

When you find alerts on the cluster overview page, especially the Critical and Warning alerts, you need to go to the cluster health status page to view and clear the alerts.

Prerequisites

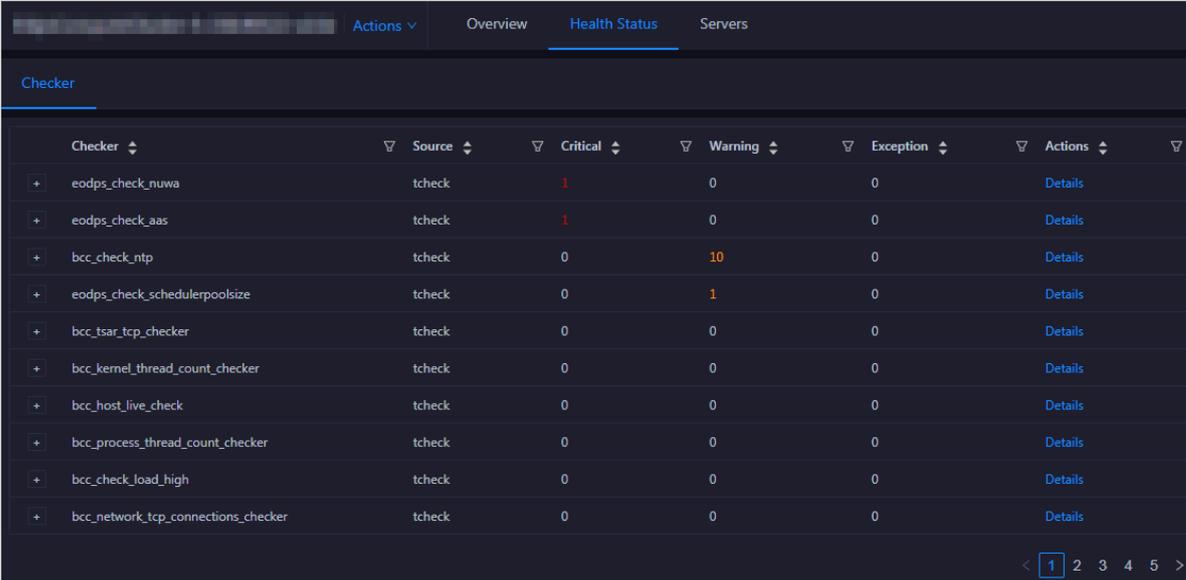
Your ABM account must have the required permissions to perform O&M operations on MaxCompute.

Context

This topic describes how to view and clear alerts for a MaxCompute cluster.

Procedure

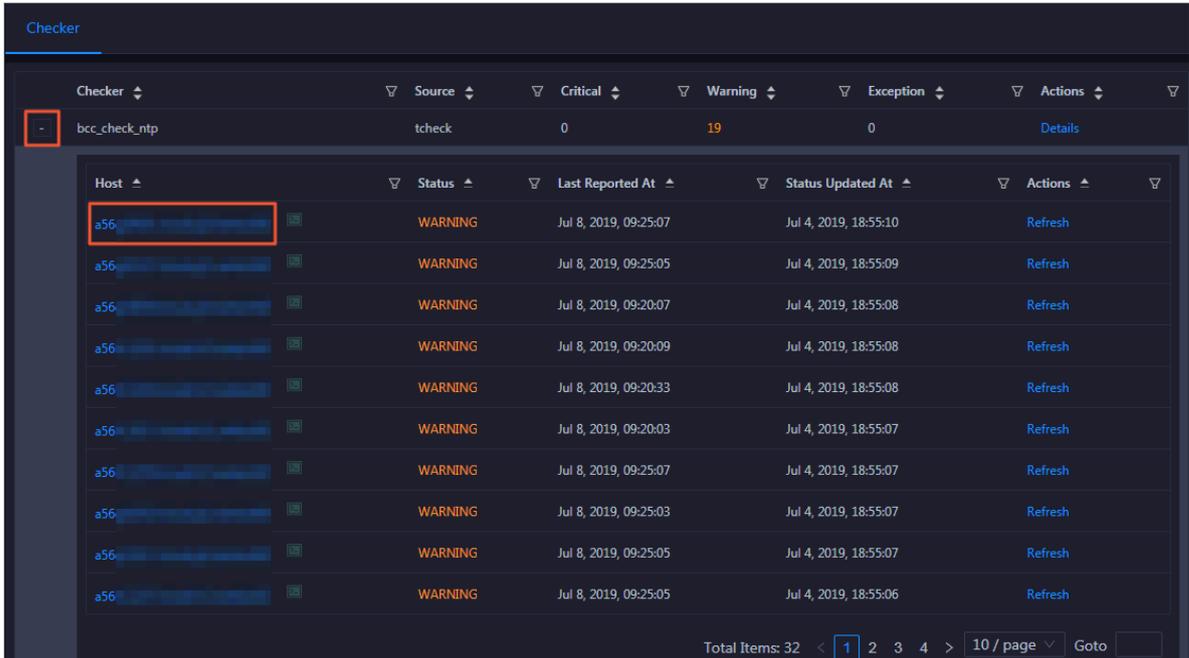
1. [Log on to the ABM console](#).
2. Click  in the upper-left corner, and select MaxCompute.
3. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.



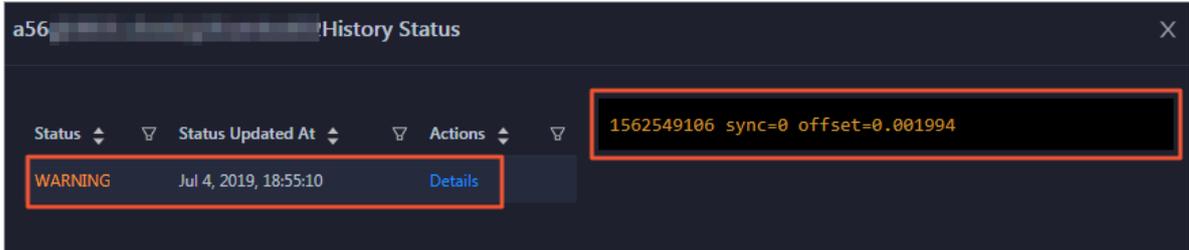
The screenshot shows the 'Health Status' tab in the MaxCompute console. It displays a table with columns for Checker, Source, Critical, Warning, Exception, and Actions. The table lists various checker items with their respective status counts.

Checker	Source	Critical	Warning	Exception	Actions
+ eodps_check_nuwa	tcheck	1	0	0	Details
+ eodps_check_aas	tcheck	1	0	0	Details
+ bcc_check_ntp	tcheck	0	10	0	Details
+ eodps_check_schedulerpoolsize	tcheck	0	1	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_host_live_check	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details
+ bcc_network_tcp_connections_checker	tcheck	0	0	0	Details

- On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.



- Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



6. On the Health Status page, click Details in the Actions column of the checker to view the schemes to clear the alerts.

Details ✕

Name:	bcc_disk_usage_checker	Source:	tcheck
Alias:	Disk Usage Check	Application:	bcc
Type:	system	Scheduling:	Enable

Data Collection: Enable

Default Execution Interval: 0 0/5 * * * ?

Description:

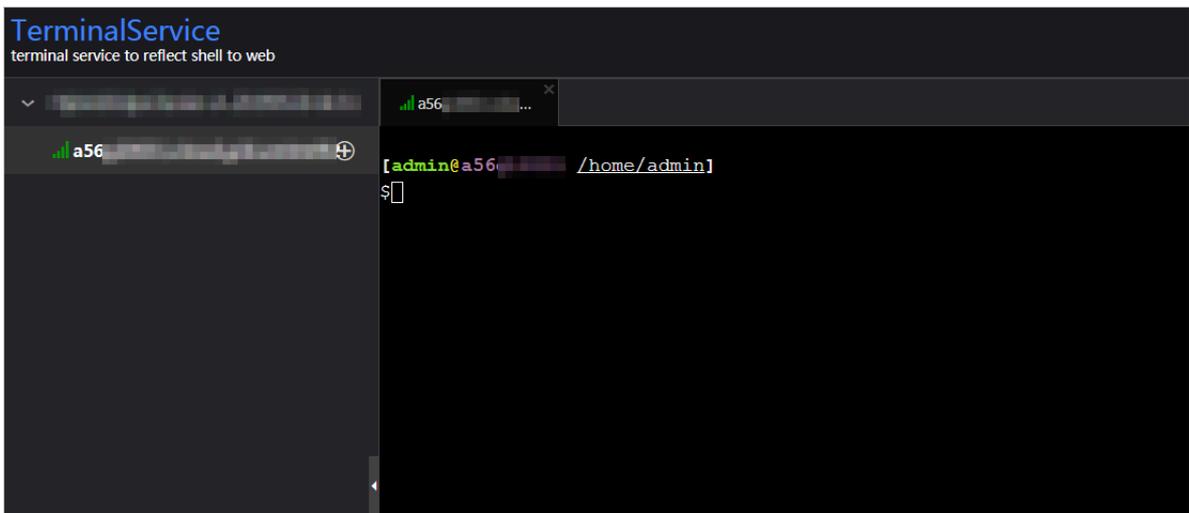
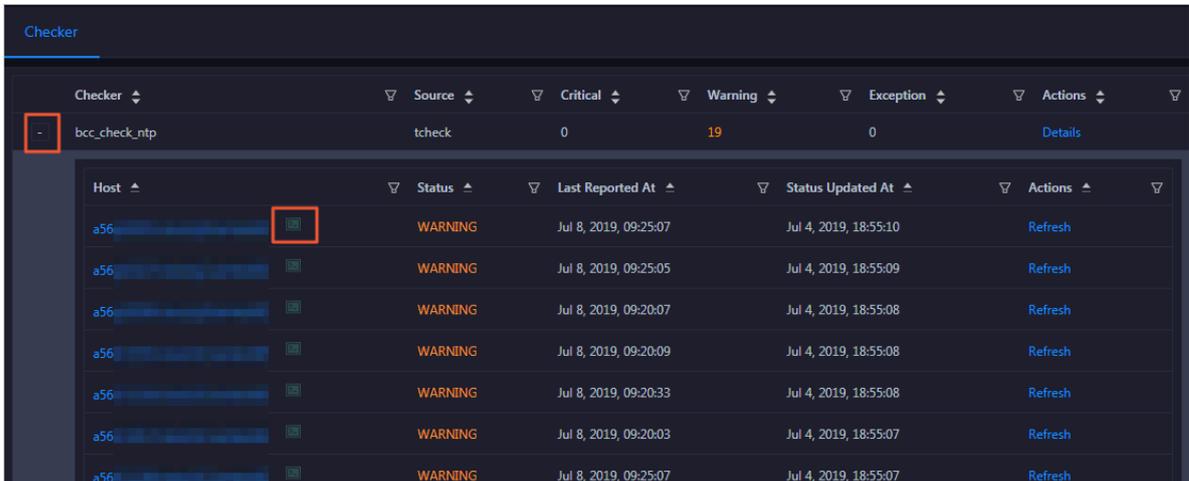
This checker checks the storage usage by using this command: `df -lh`. A warning is triggered when the usage exceeds 80% and a critical alert is triggered when the usage exceeds 90%. Reason: User operations. Old log data is not deleted. Logrotate is not working. Fix:

1. Log on to the server and list all partitions by executing this command: `df -lh`
2. Execute the following command on each partition to find the directory where the error occurred: `du -sh *`
3. Determine the cause of the issue and find a solution. You can create a task to clear log data periodically.

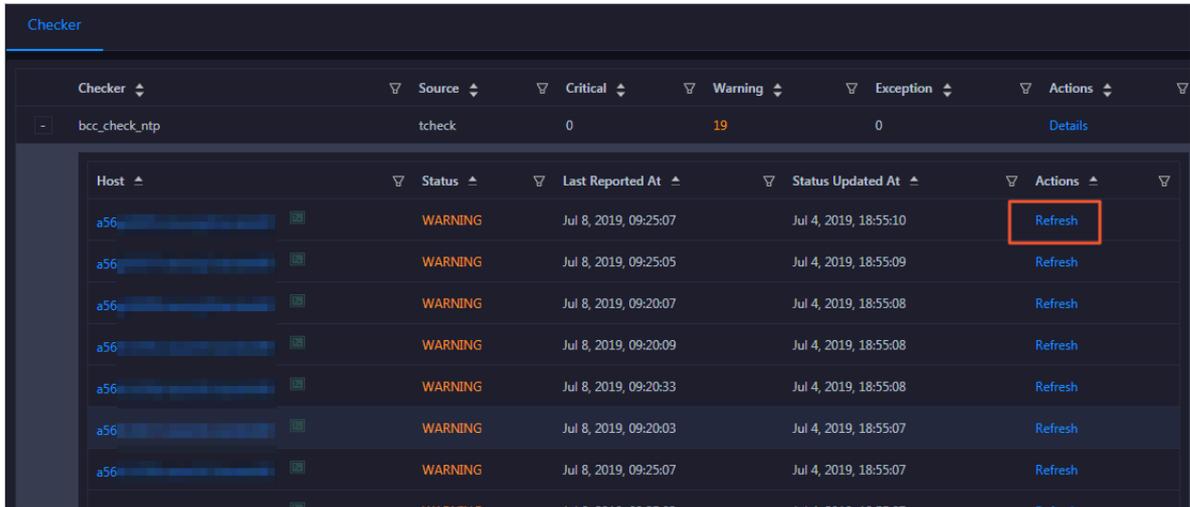
> Show More

7. Clear the alerts according to the schemes.

To log on to a host with alerts for related operations, click the Log On icon next to the name of the host. On the TerminalService page that appears, click the hostname on the left to log on to the host.



8. After you clear an alert for a host, click **Refresh** in the **Actions** column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



The screenshot shows a 'Checker' dashboard with a table of hosts. The table has columns for Host, Status, Last Reported At, Status Updated At, and Actions. The first row shows a host with a 'WARNING' status and a 'Refresh' button highlighted with a red box. The table also shows summary statistics at the top: Critical: 0, Warning: 19, Exception: 0.

Checker	Source	Critical	Warning	Exception	Actions
bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

3.1.4 ABM

3.1.4.1 ABM dashboard

The Apsara Bigdata Manager (ABM) dashboard displays key operation metrics for the four core services, namely MaxCompute, DataWorks, Realtime Compute, and DataHub. It also provides information about alerts for all big data services, so that you can understand the overall running status of the big data services. In addition, the dashboard supports automatic data refresh and full-screen display.

Entry

Log on to the ABM console. The Dashboard page appears. To return to the Dashboard page from any other page, click  in the upper-left corner and then click ABM.

The screenshot shows the Apsara Bigdata Manager interface. At the top, there's a navigation bar with 'Apsara Bigdata Manager | ABM' and tabs for 'Monitor', 'O&M', and 'Management'. Below this is a 'Dashboard' section with a region dropdown set to 'cn' and a time range from 'Jul 8, 2019, 15:31:49' to 'Jul 8, 2019, 17:31:49'. The main content is an 'Overview' table with columns for 'Products', 'OK', 'Critical', 'Warning', and 'Exception'. Below the table is a pagination bar showing 'Total Items: 9' and '50 / page'. A sidebar on the left lists expandable categories: MaxCompute, DataWorks, StreamCompute, and DataHub.

Products	OK	Critical	Warning	Exception
MaxCompute	1319	6	63	0
blink	324	0	38	0
DataHub	1034	0	16	0
ElasticSearch	84	24	0	0
BigGraph	762	0	58	0
ABM	108	4	0	0
QuickBI	112	0	0	0
PAI	128	0	0	0
I+	96	0	0	0

On the Dashboard page, you can select a region from the region drop-down list in the upper-left corner. In this way, you can view the cluster running status of each big data service in the specified region.

View and clear alerts of various services

In the overview section, you can view the number of alerts for each of all big data services. You need to pay special attention to the Critical and Warning alerts. These alerts must be cleared in a timely manner.

1. On the Dashboard page, click the number of Critical or Warning alerts of a service such as MaxCompute in the overview section. The Health Status page under Clusters of the service appears.

Checker	Source	Critical	Warning	Exception	Actions
+ eodps_check_nuwa	tcheck	1	0	0	Details
+ eodps_check_aas	tcheck	1	0	0	Details
+ bcc_check_ntp	tcheck	0	10	0	Details
+ eodps_check_schedulerpoolsize	tcheck	0	1	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_host_live_check	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details
+ bcc_network_tcp_connections_checker	tcheck	0	0	0	Details

On the Health Status page, you can view all the checkers of the service.

2. Click Details in the Actions column of a checker with alerts. In the Details dialog box that appears, view the details of the checker and the scheme to clear the alerts. Follow the steps in the scheme to clear the alerts.

Name: eodps_check_nuwa **Source:** tcheck

Alias: Nuwa Check **Application:** eodps

Type: system **Scheduling:** Enable

Data Collection: Enable

Default Execution Interval: 0 0/30 * * * ?

Description:

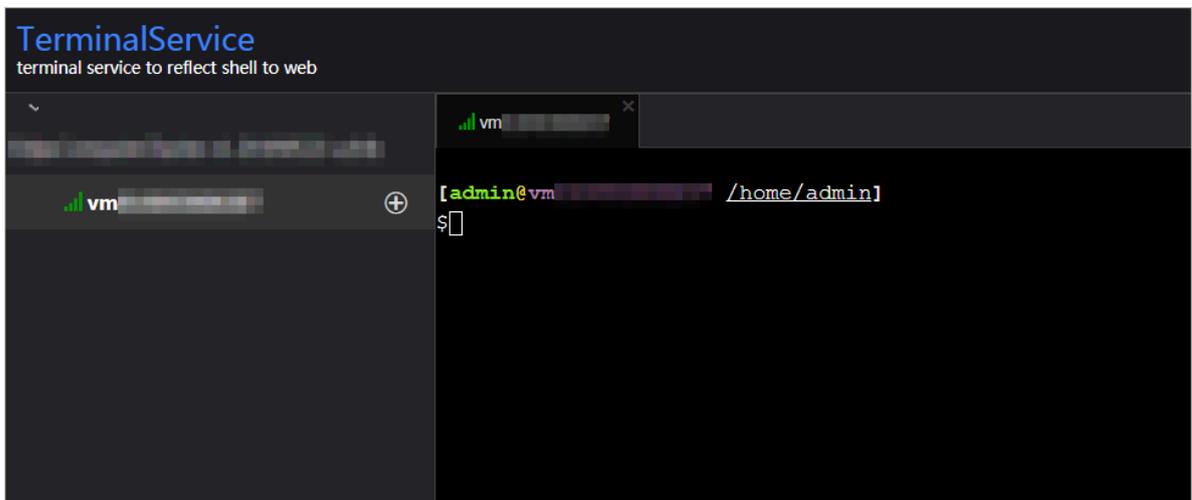
Fix:

1. Check the status of Nuwa services: `echo svr|nc localhost 10240|grep Mode|grep follower;echo svr|nc localhost 10240|grep Mode|grep leader;`
2. Submit a ticket to report the result.

> Show More

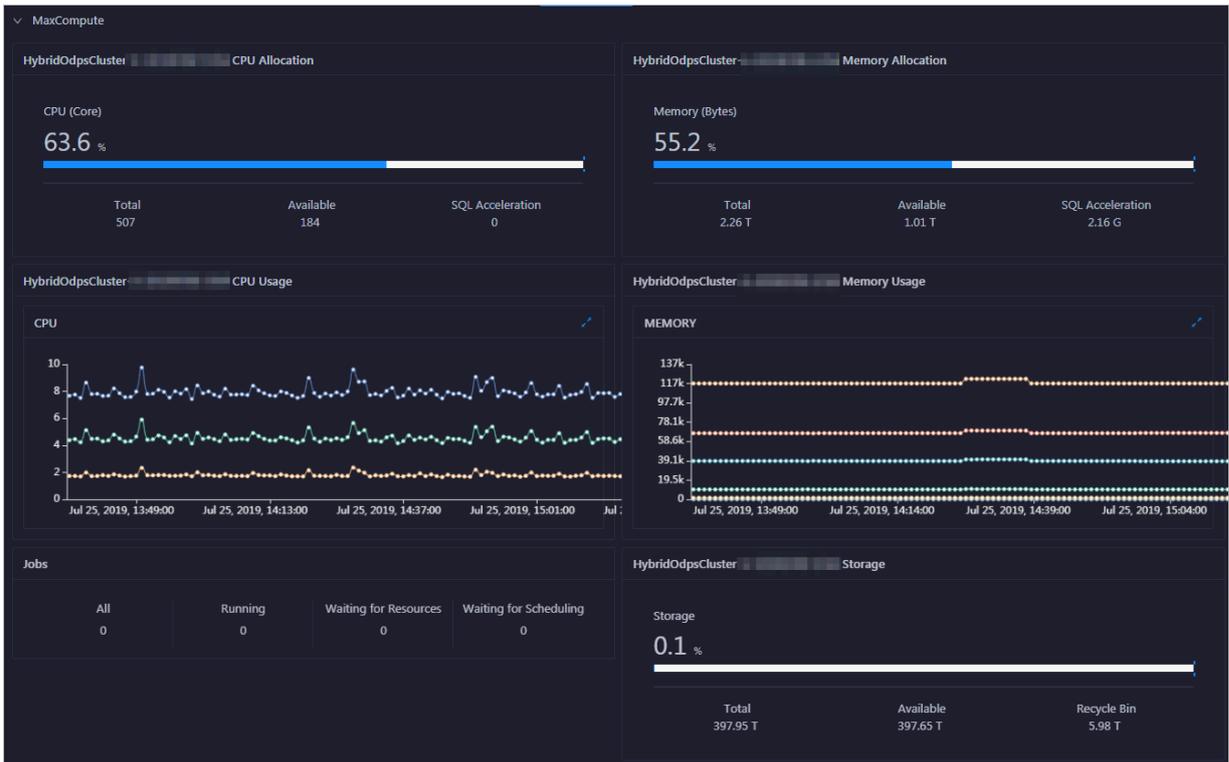
3. Log on to the hosts with alerts if necessary for related operations.

Click + to expand a checker with alerts, and then click the Log On icon next to the name of a host with alerts. On the TerminalService page that appears, click the hostname on the left to log on to the host.



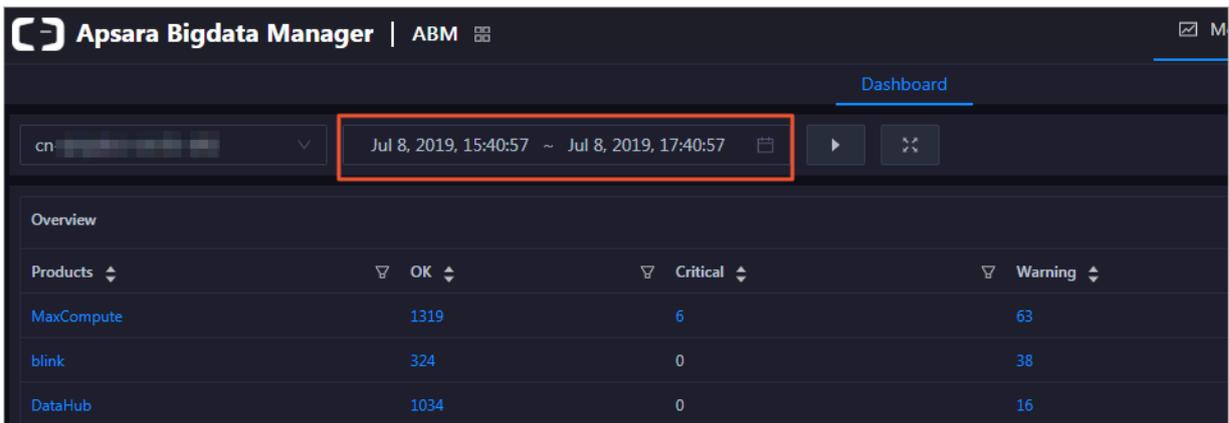
View key operation metrics for MaxCompute

The ABM dashboard displays key operation metrics for MaxCompute. On the Dashboard page, click MaxCompute in the overview section to view the metrics.



In the MaxCompute section, you can view the allocation of CPU and memory, CPU and memory usage trends, job running status, and storage usage of the MaxCompute cluster. For more information about the MaxCompute operation metrics, see [Cluster overview](#).

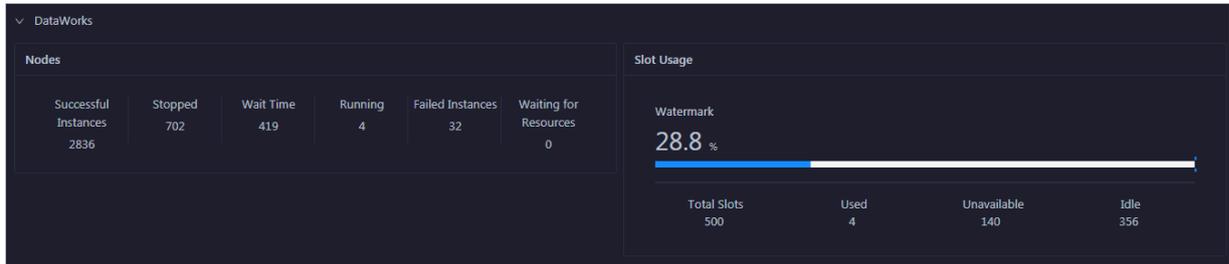
If automatic refresh is not enabled for the trend charts of CPU and memory usage, you can specify a time period at the top of the Dashboard page to view the CPU and memory usage in this period.



If automatic refresh is enabled, the trend charts of CPU and memory usage are displayed based on the refresh settings specified for automatic refresh.

View key operation metrics for DataWorks

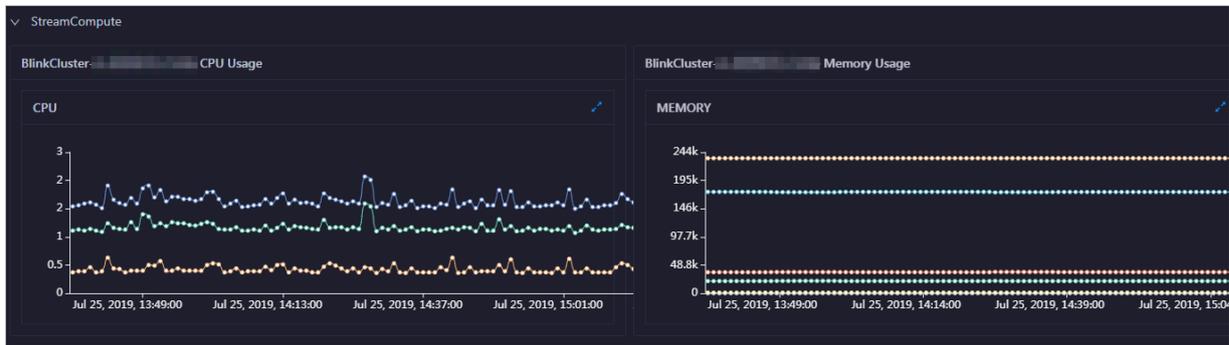
The ABM dashboard displays key operation metrics for DataWorks. On the Dashboard page, click DataWorks in the overview section to view the metrics.



In the DataWorks section, you can view the node scheduling and slot usage of the DataWorks cluster. For more information about the DataWorks operation metrics, see [DataWorks overview](#).

View key operation metrics for Realtime Compute

The ABM dashboard displays key operation metrics for Realtime Compute. On the Dashboard page, click StreamCompute in the overview section to view the metrics.



In the StreamCompute section, you can view the CPU and memory usage trend charts of the Realtime Compute cluster. For more information about the Realtime Compute operation metrics, see [Cluster overview](#).

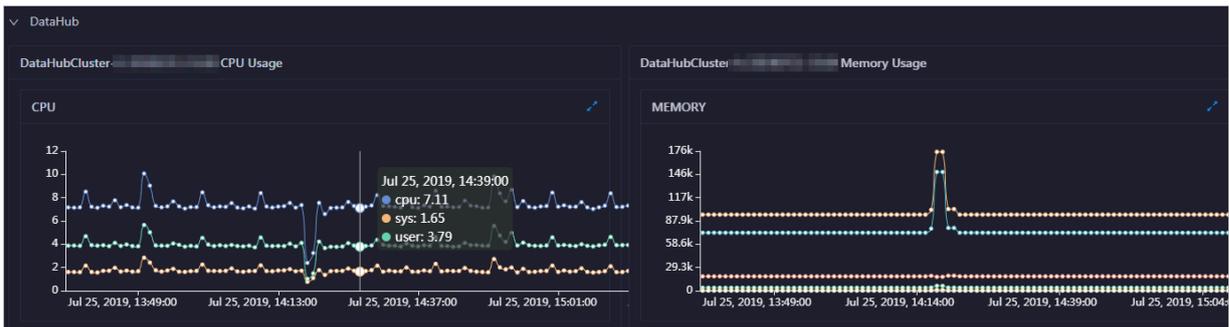
If automatic refresh is not enabled for the trend charts of CPU and memory usage, you can specify a time period at the top of the Dashboard page to view the CPU and memory usage in this period.



If automatic refresh is enabled, the trend charts of CPU and memory usage are displayed based on the refresh settings specified for automatic refresh.

View key operation metrics for DataHub

The ABM dashboard displays key operation metrics for DataHub. On the Dashboard page, click DataHub in the overview section to view the metrics.



In the DataHub section, you can view the CPU and memory usage trend charts of the DataHub cluster. For more information about the DataHub operation metrics, see [Cluster overview](#).

If automatic refresh is not enabled for the trend charts of CPU and memory usage, you can specify a time period at the top of the Dashboard page to view the CPU and memory usage in this period.

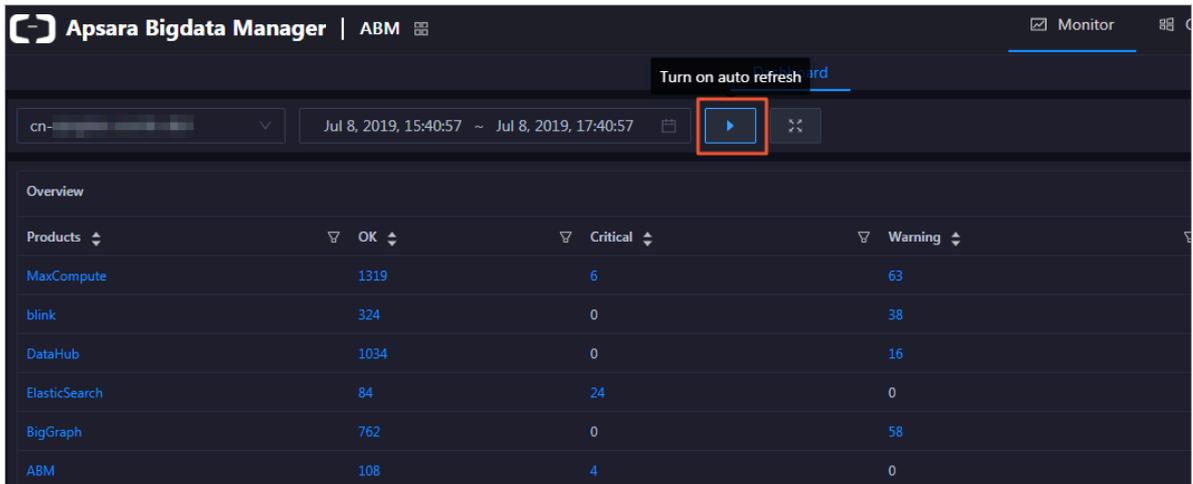


If automatic refresh is enabled, the trend charts of CPU and memory usage are displayed based on the refresh settings specified for automatic refresh.

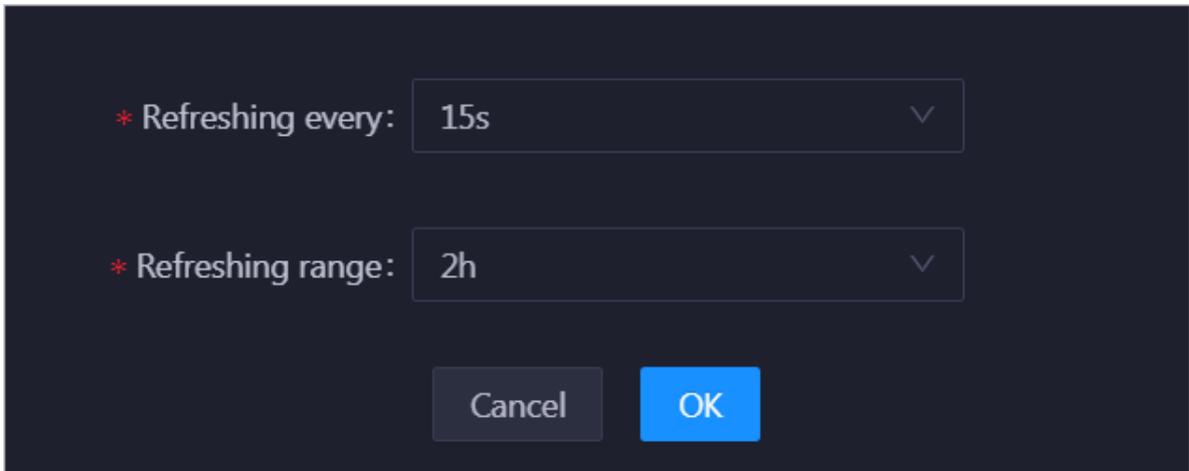
Enable or disable automatic refresh

By default, automatic refresh is disabled for the metrics on the dashboard. You can enable the feature as required.

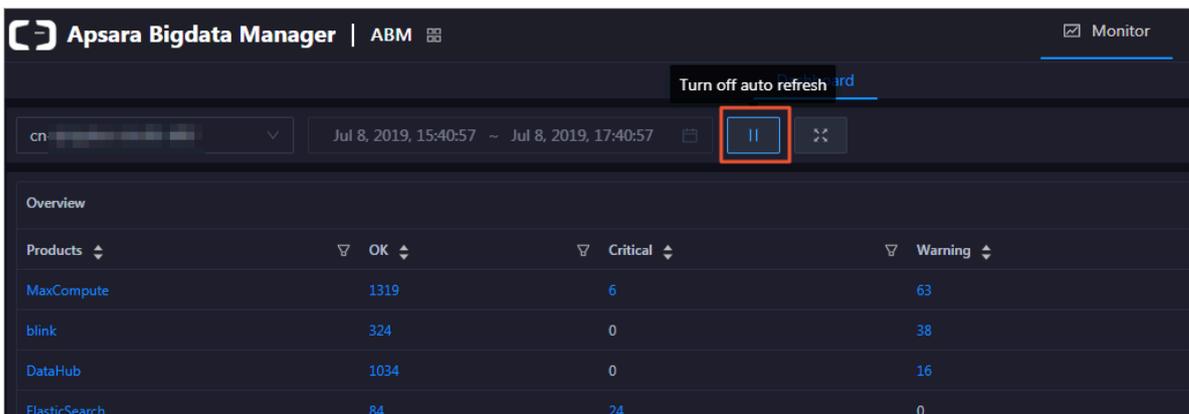
1. At the top of the Dashboard page, click the Turn on auto refresh icon.



2. In the dialog box that appears, select the refresh interval and time range. The Refreshing range field specifies the time range of the trend charts of CPU and memory usage for each cluster.



3. Click OK.



When automatic refresh is enabled, the Turn on auto refresh icon is replaced with the Turn off auto refresh icon. The system automatically refreshes all data on the dashboard according to the specified time interval.

To disable automatic refresh, click the Turn off auto refresh icon.

Display the dashboard in full-screen mode

The dashboard provides a full-screen display feature for you to view the running status of big data services clearly.

At the top of the Dashboard page, click the icon for full-screen display to display the Dashboard page in full-screen mode.

Products	OK	Critical	Warning
MaxCompute	1319	6	63
blink	324	0	38
DataHub	1034	0	16
ElasticSearch	84	24	0
BigGraph	762	0	58

3.1.4.2 ABM repository

This topic describes the features of the ABM repository and how to access the Repository page.

Modules

- For MaxCompute, the Repository page displays the trend charts of CU and storage usage, records of CU and storage usage, and proportions of idle CUs and storage.
- For DataWorks, the Repository page displays the trend chart of slot usage, records of slot usage, and proportion of idle slots.
- For DataHub, the Repository page displays the trend chart of storage usage, storage usage details, and the percent of idle storage.

Entry

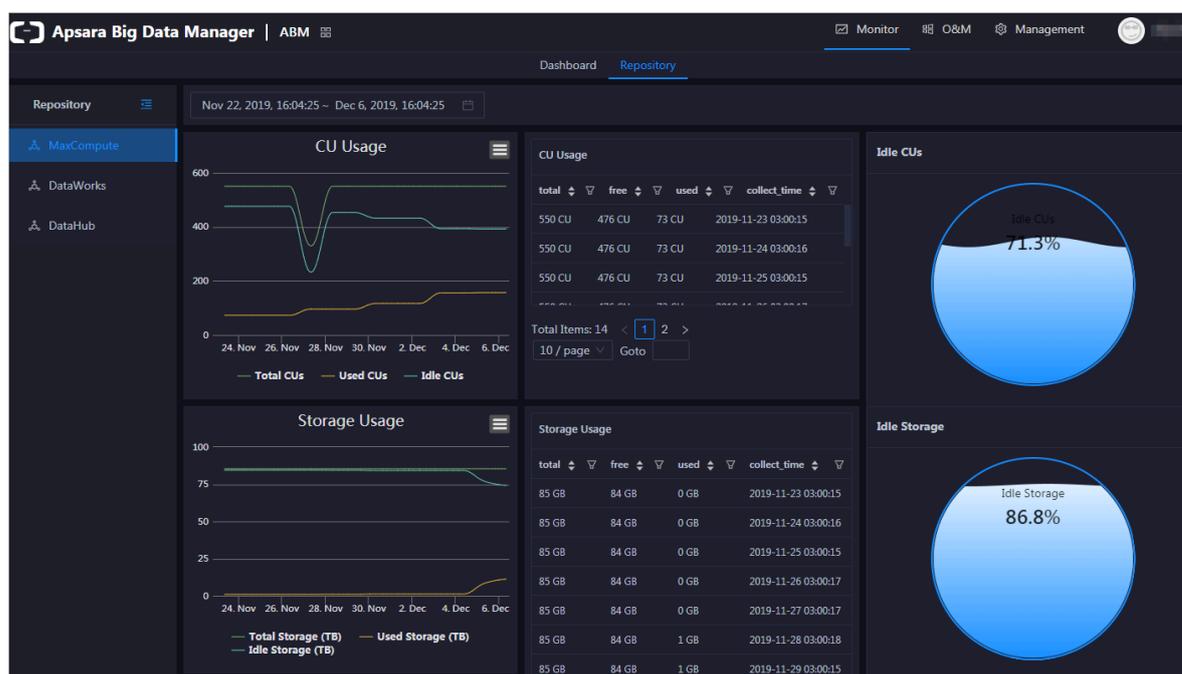
1. *Log on to the ABM console.* The Dashboard page appears.



Note:

To return to the Dashboard page from any other page, click  in the upper-left corner and then click ABM.

2. Click the Repository tab. The Repository page appears.



Supported operations

You can filter or sort records of CU, storage, and slot usage based on a column to facilitate information retrieval. For more information, see [Common operations](#).

3.1.4.3 O&M overview

This topic describes the features of Apsara Bigdata Manager (ABM) O&M and how to access the ABM O&M page.

Modules

ABM O&M includes service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Service O&M	Service overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission , TCP connection, and root disk usage for each service in a cluster.
	Service hosts	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.

Module	Feature	Description
Cluster O&M	Cluster overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Cluster health	Displays the check results for a cluster. The check results are divided into the Critical, Warning, Exception, and OK types.
Host O&M	Host overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Host health	Displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click ABM.

3. On the ABM page, click O&M in the upper-right corner. By default, the Services page appears.



The O&M page consists of the Services, Clusters, and Hosts modules.

3.1.4.4 Cluster O&M

3.1.4.4.1 Cluster overview

The cluster overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.

Entry

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.



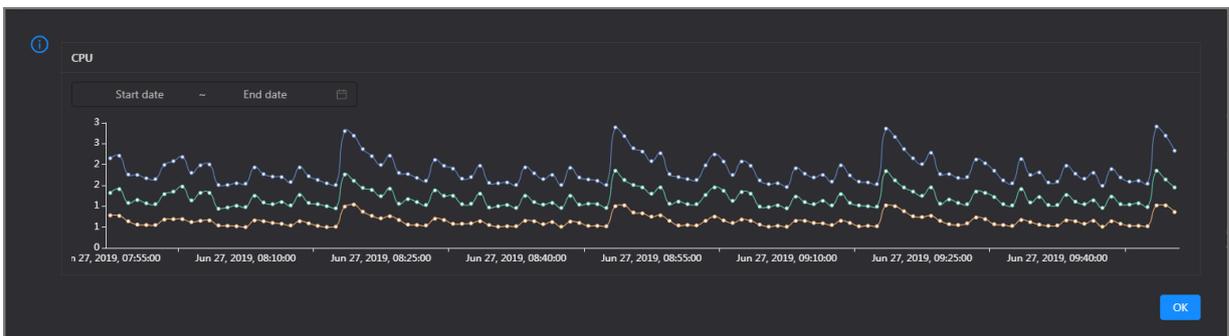
The cluster overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster. The trend charts are described as follows:

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

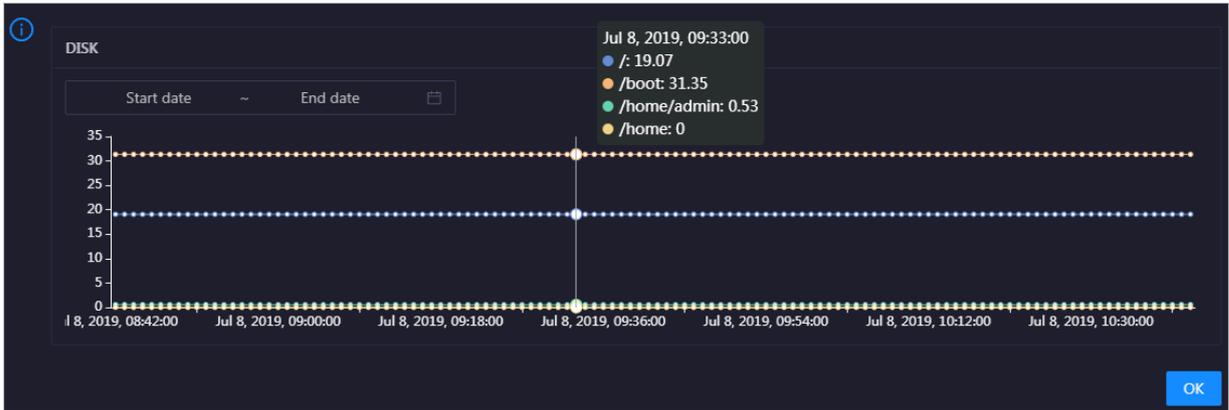
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the /, /boot, /home/admin, and /home directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

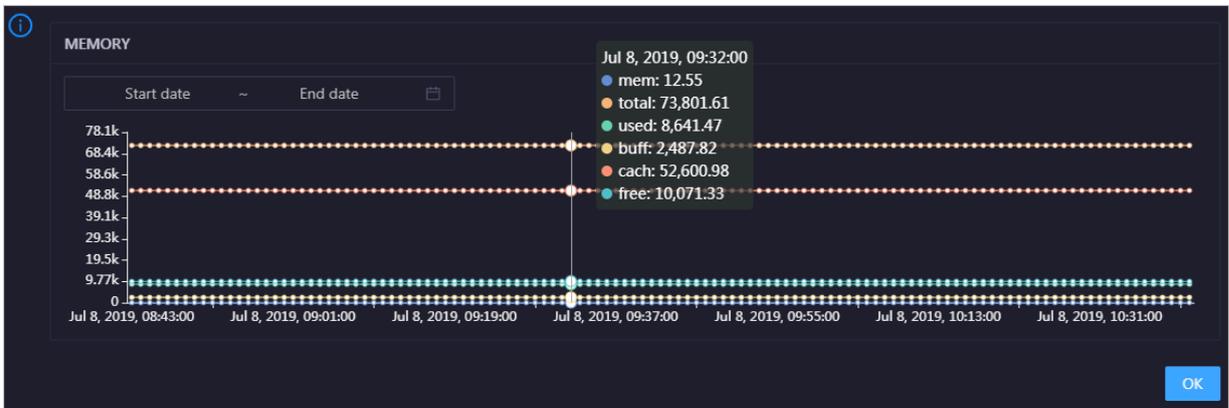


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

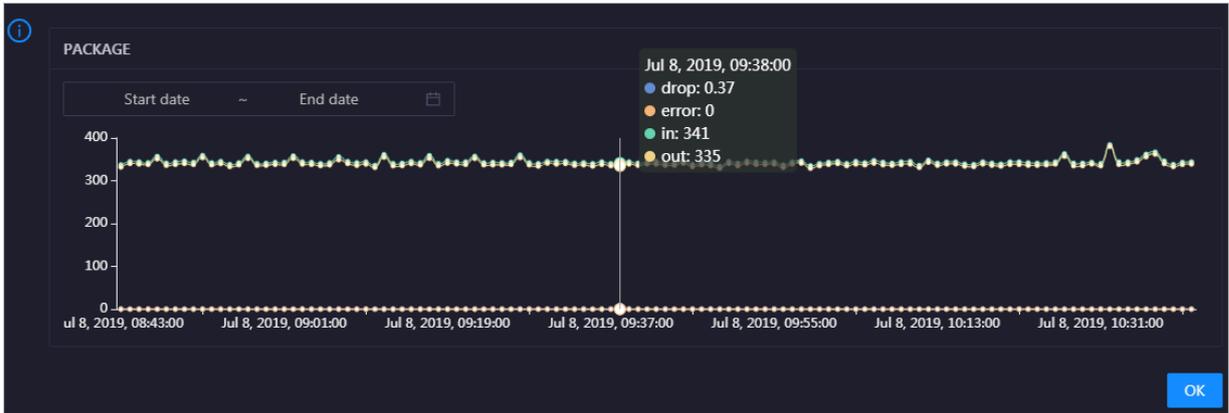


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

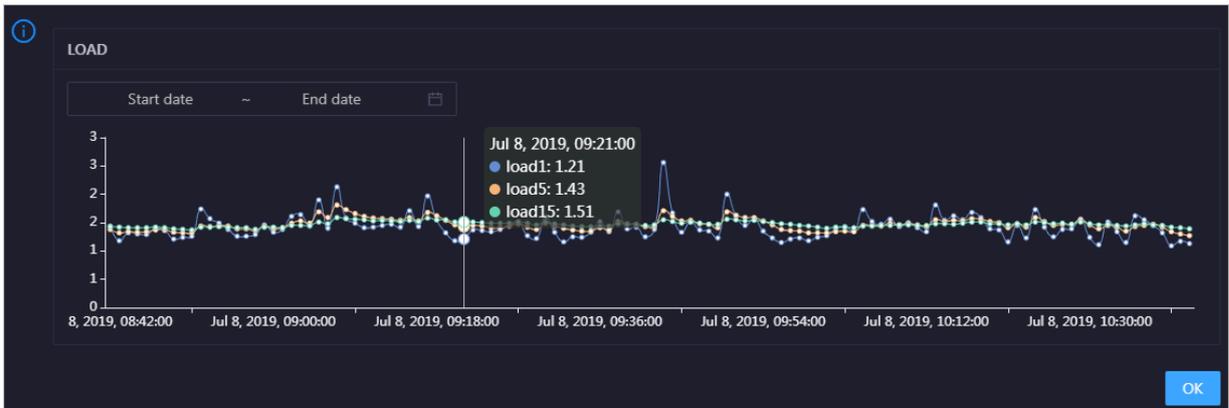


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

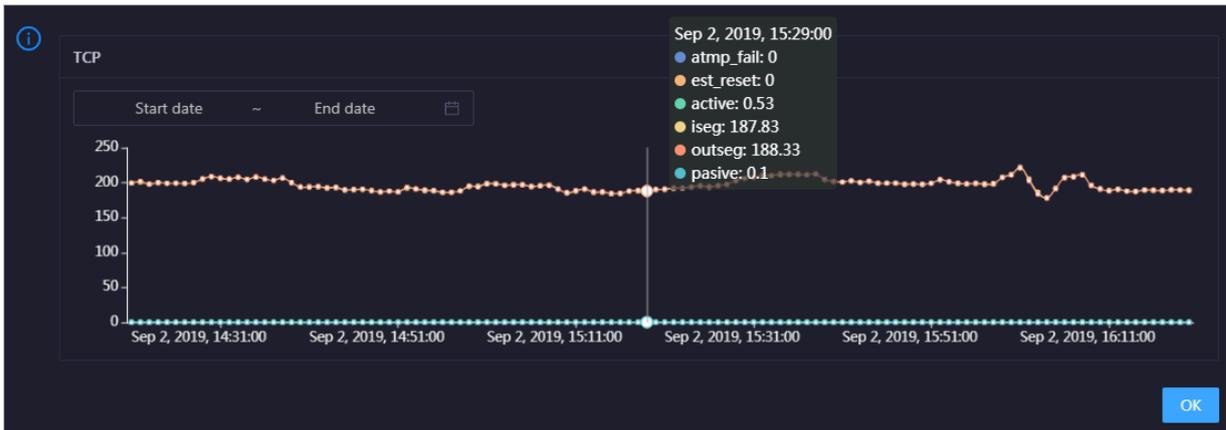


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

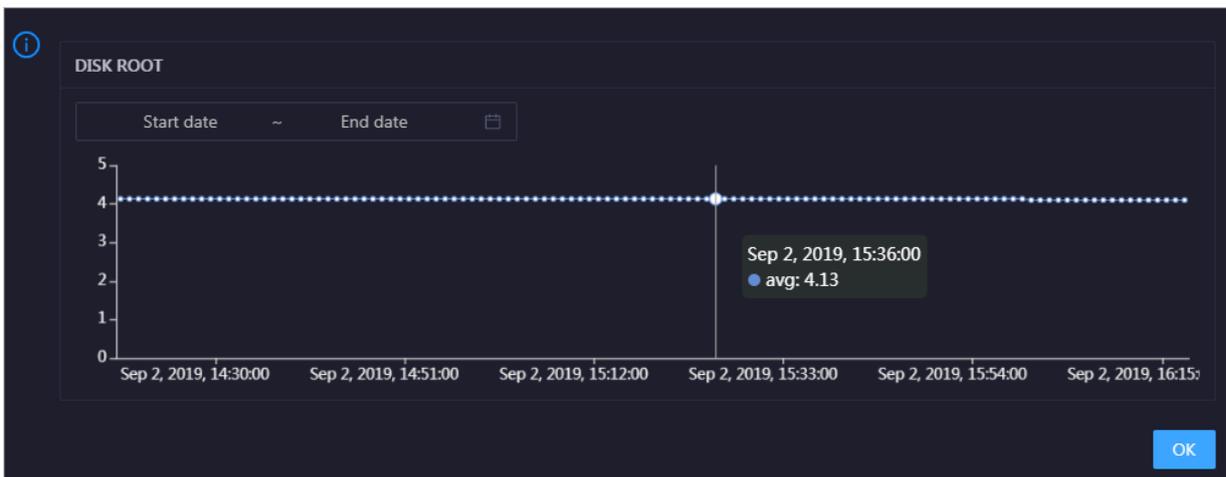


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

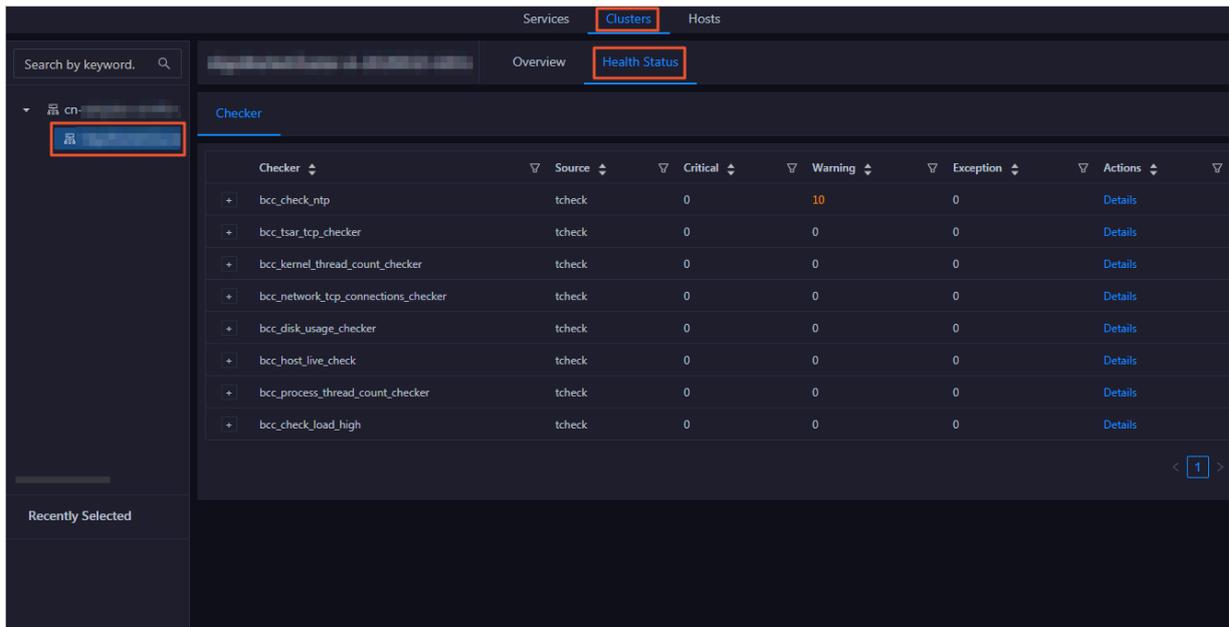
3.1.4.4.2 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for hosts in the cluster, and schemes to clear

alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

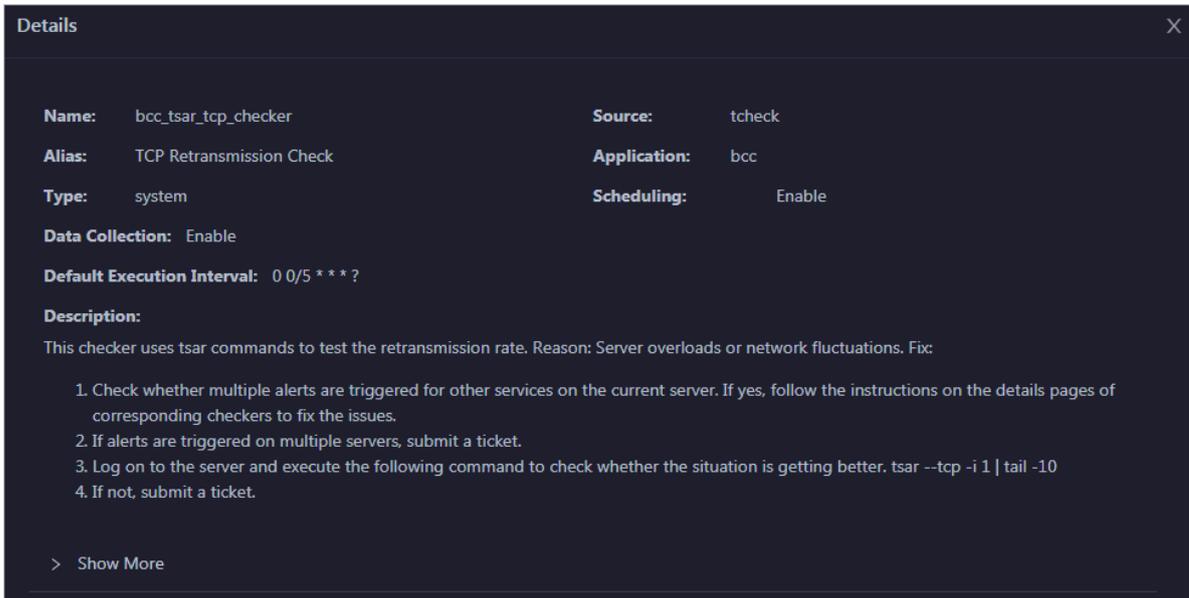
On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.



On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception events are alerts. You need to pay attention to them, especially the Critical and Warning events.

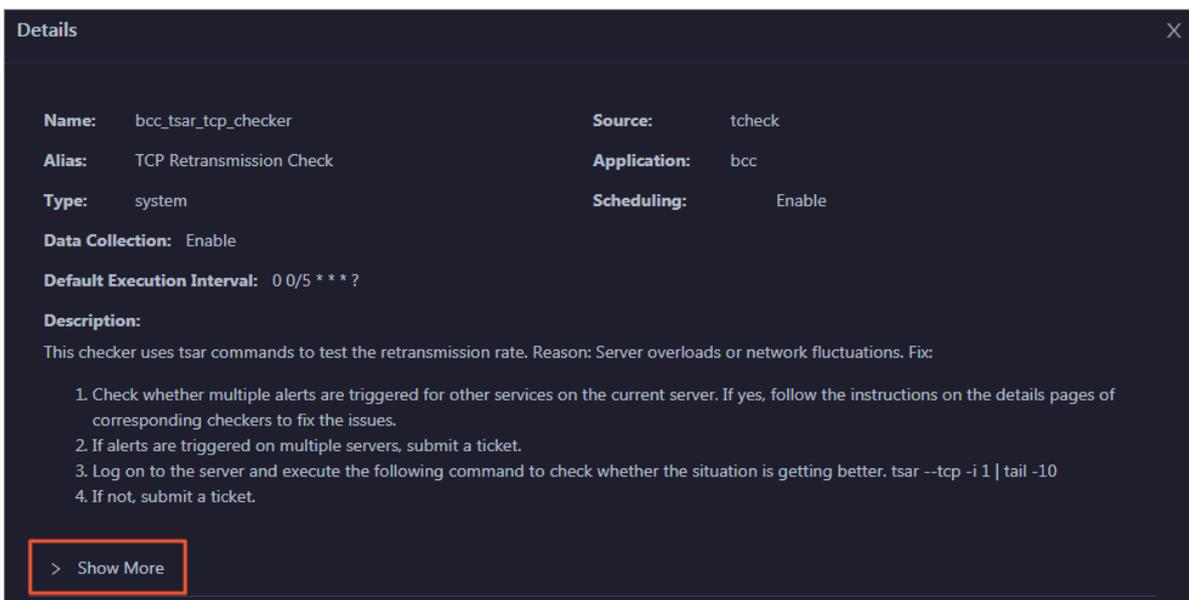
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

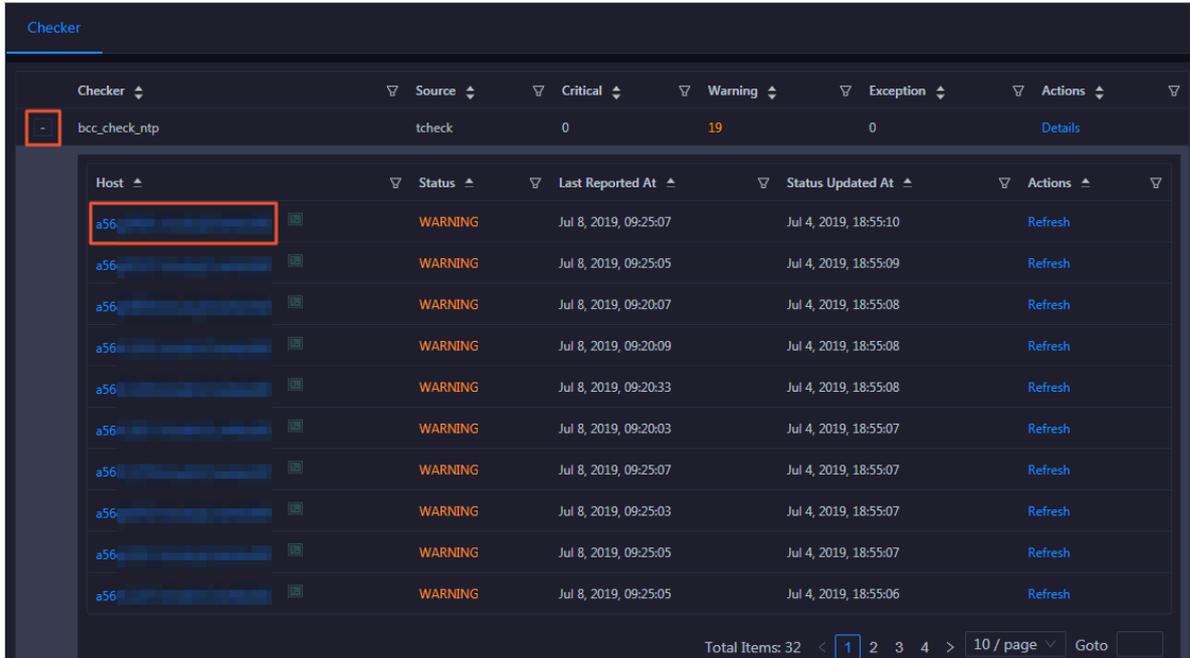


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

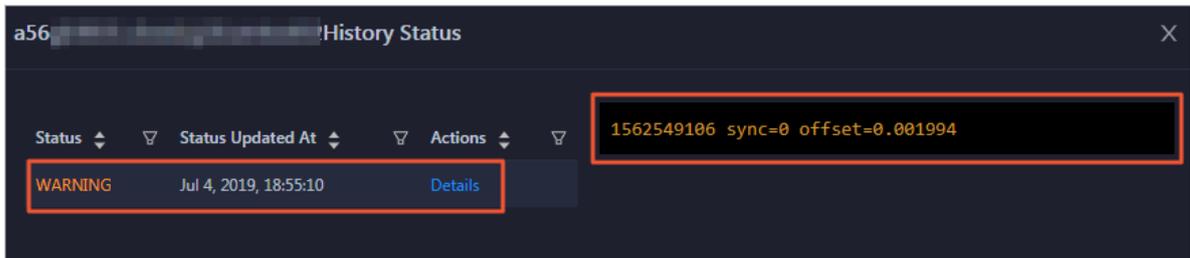
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

- 1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.**



- 2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**



Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

Details ✕

Name: bcc_disk_usage_checker	Source: tcheck
Alias: Disk Usage Check	Application: bcc
Type: system	Scheduling: Enable
Data Collection: Enable	
Default Execution Interval: 0 0/5 * * * ?	

Description:
 This checker checks the storage usage by using this command: `df -lh`. A warning is triggered when the usage exceeds 80% and a critical alert is triggered when the usage exceeds 90%. Reason: User operations. Old log data is not deleted. Logrotate is not working. Fix:

1. Log on to the server and list all partitions by executing this command: `df -lh`
2. Execute the following command on each partition to find the directory where the error occurred: `du -sh *`
3. Determine the cause of the issue and find a solution. You can create a task to clear log data periodically.

> Show More

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

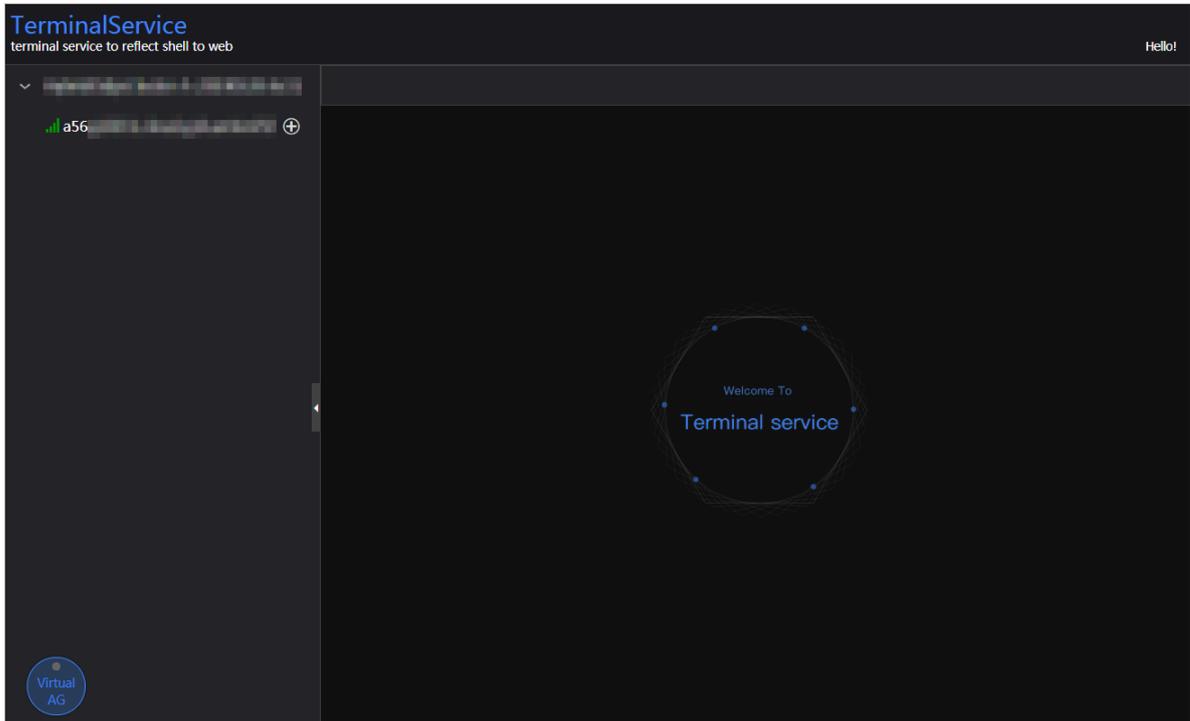
- 1. On the Health Status page, click + to expand a checker with alerts.**

Checker

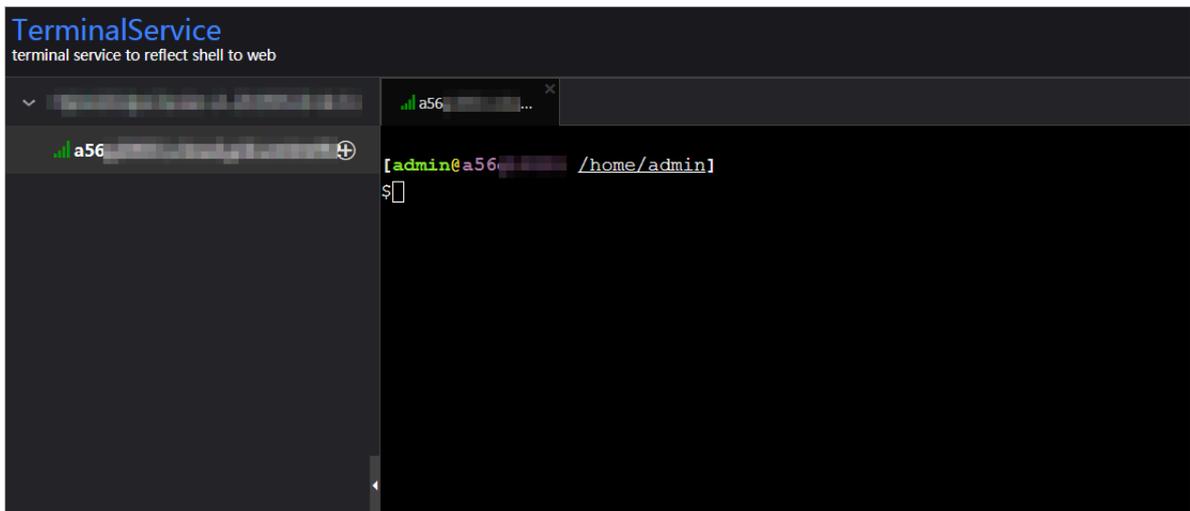
Checker	Source	Critical	Warning	Exception	Actions
- bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56 +	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56 +	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56 +	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.



3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

Checker	Source	Critical	Warning	Exception	Actions
bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

3.1.4.5 Service O&M

3.1.4.5.1 Service overview

The service overview page lists all Apsara Bigdata Manager (ABM) services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.

Entry

On the Services page, select a cluster above the left-side service list, select a service in the service list, and then click the Overview tab. The Overview page for the service appears.



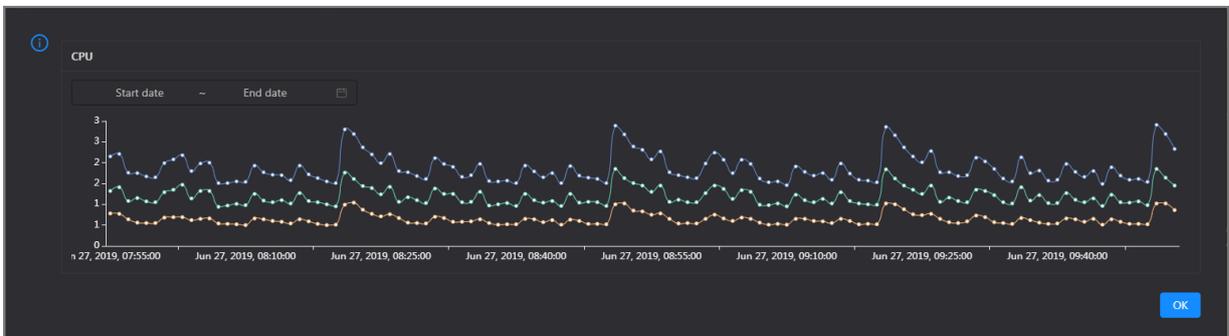
On the Overview page, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

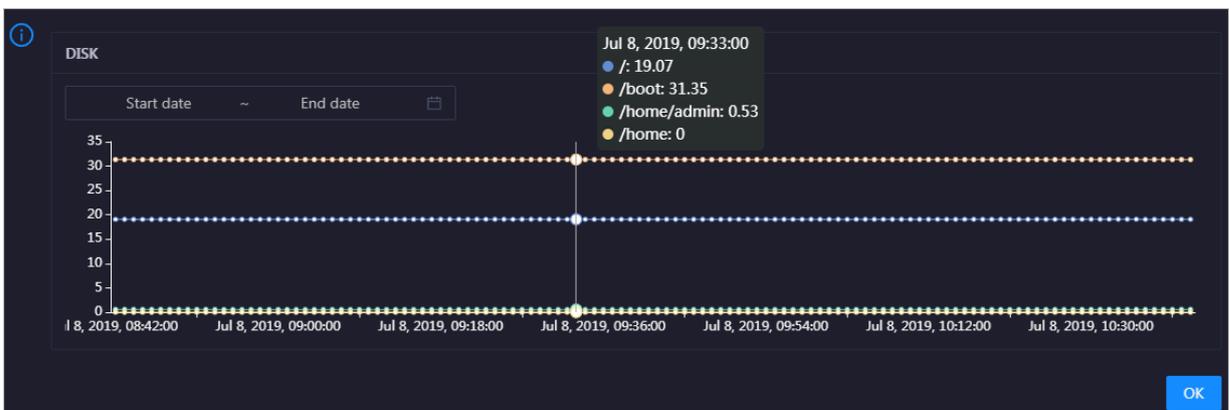
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the service in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the /, /boot, /home/admin, and /home directories for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

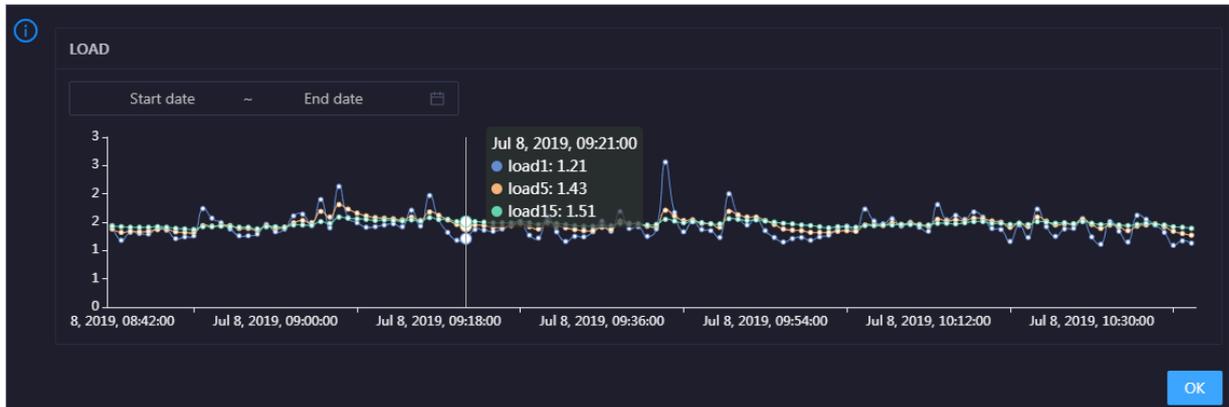


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

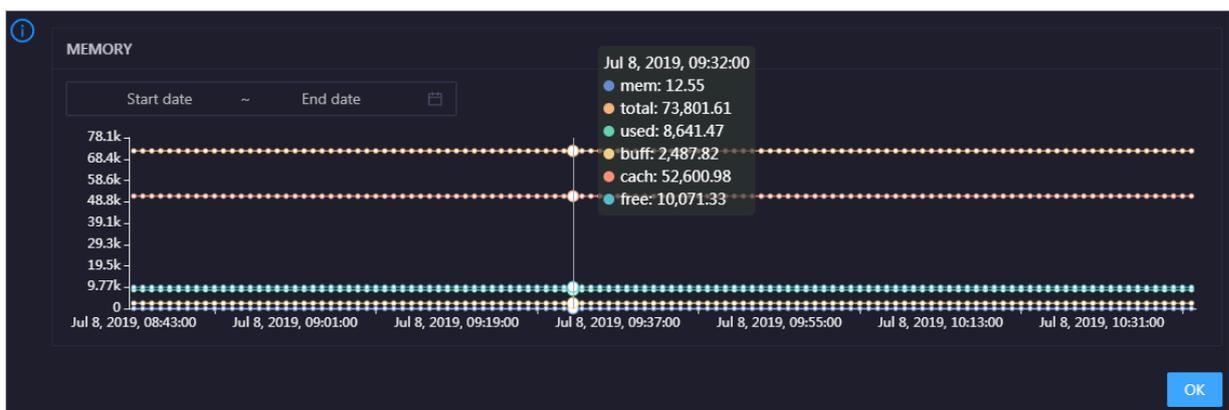


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

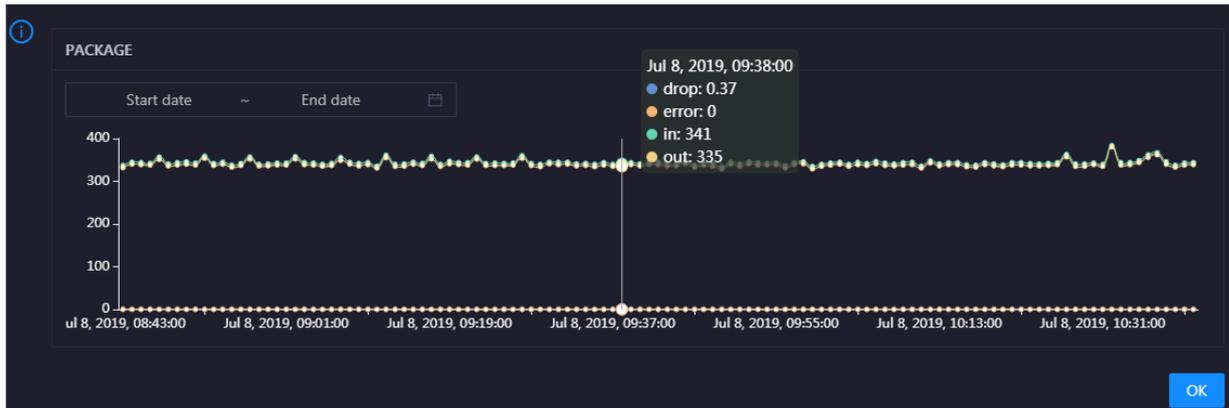


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the selected service over time in different colors. These trend lines reflect the data transmission status of the service.

Click  in the upper-right corner of the chart to zoom in it.

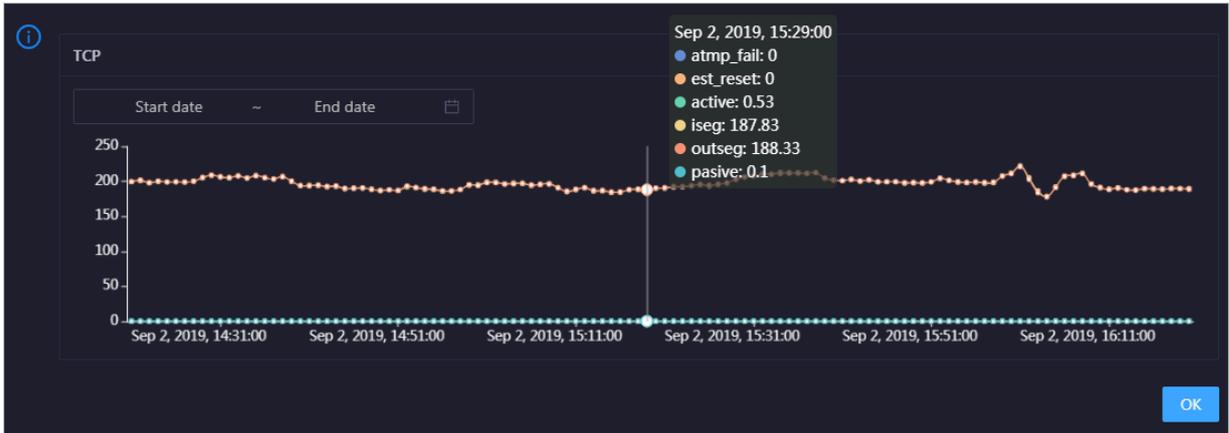


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the selected service over time in different colors. These trend lines reflect the TCP connection status of the service.

Click  in the upper-right corner of the chart to zoom in it.

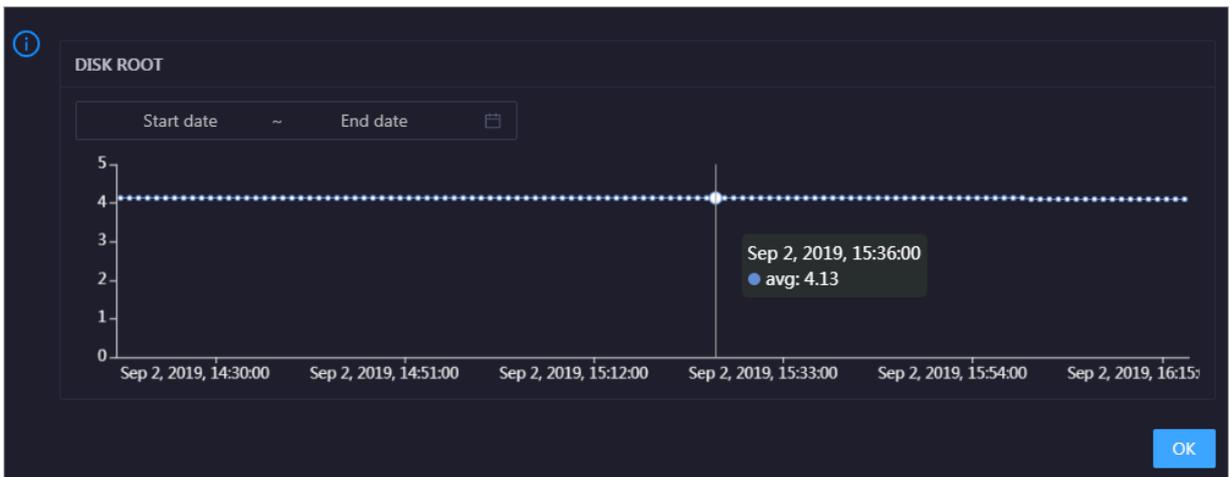


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

This chart displays the trend line of the average root disk usage (avg) for the selected service over time.

Click  in the upper-right corner of the chart to zoom in it.

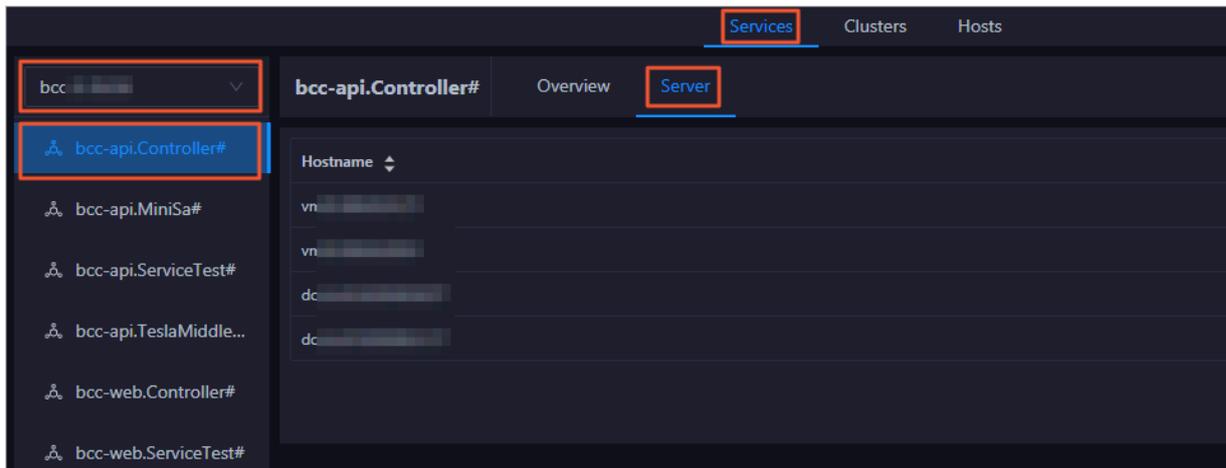


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

3.1.4.5.2 Service hosts

Apsara Bigdata Manager (ABM) allows you to view the host list of each ABM service so that you can understand the service deployment on hosts.

On the Services page, select a cluster above the left-side service list, select a service in the service list, and then click the Server tab. The Server page of the service appears.



On the Server page, you can view the hosts where the selected service is run.

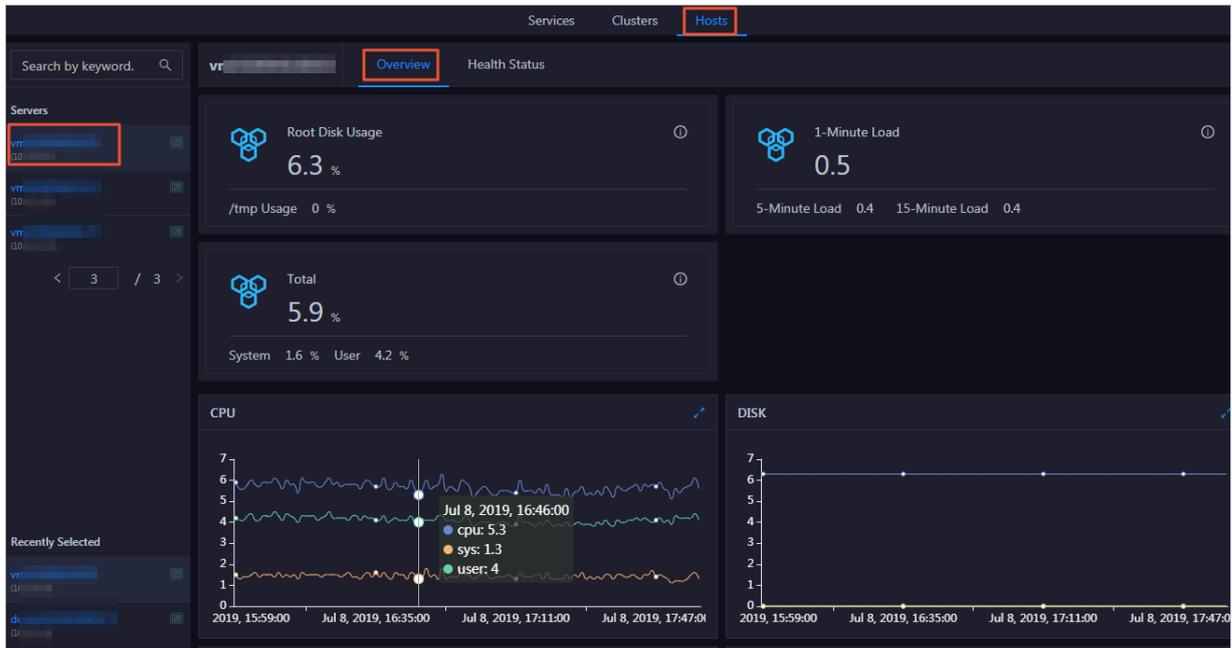
3.1.4.6 Host O&M

3.1.4.6.1 Host overview

The host overview page displays the overall running information about a host in an ABM cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Entry

On the Hosts page, select a host in the left-side navigation pane. The Overview page of the host appears.

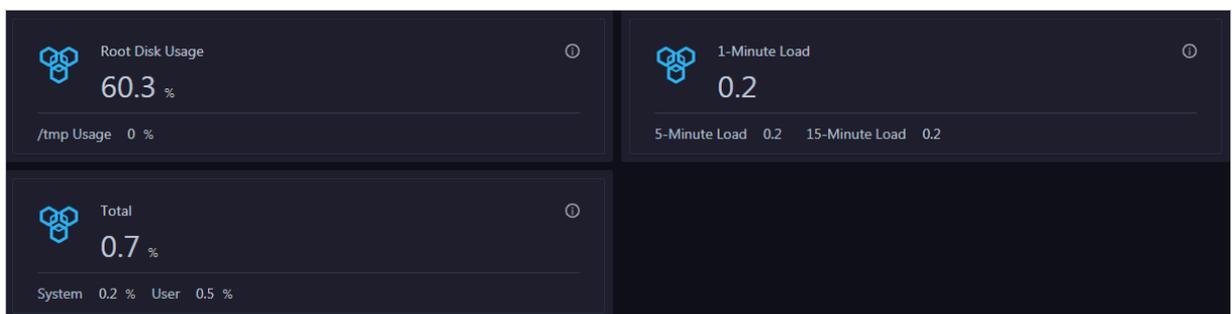


The Overview page consists of the following areas:

- **Left-side navigation pane:** displays a navigation tree of hosts.
- **Recently Selected section:** displays the recently selected hosts, which allows you to quickly switch between commonly used hosts.
- **Right pane:** displays the root disk usage, total usage, load, health check result, health check history of the host. It also displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Root Disk Usage, Total, and 1-Minute Load

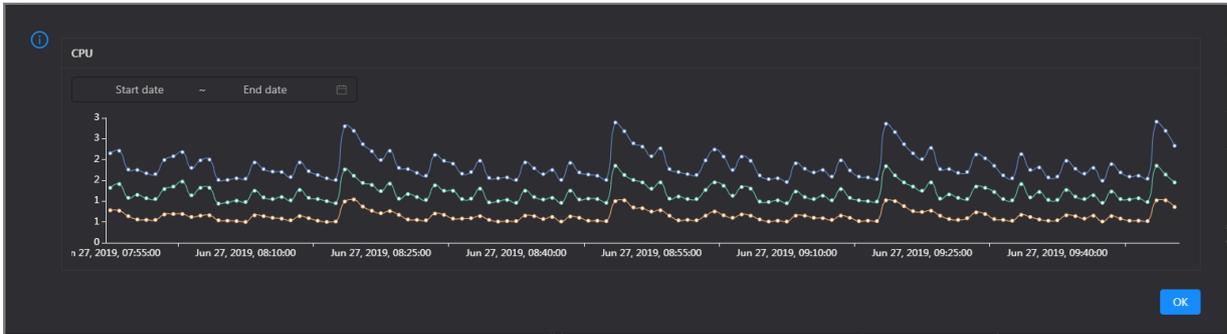
These sections display the root disk usage, total usage, and 1-minute load for the selected host. The Root Disk Usage section provides the usage of the */tmp* directory. The Total section provides the system usage and user usage. The 1-Minute Load section provides the 1-minute, 5-minute, and 15-minute load averages.



CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

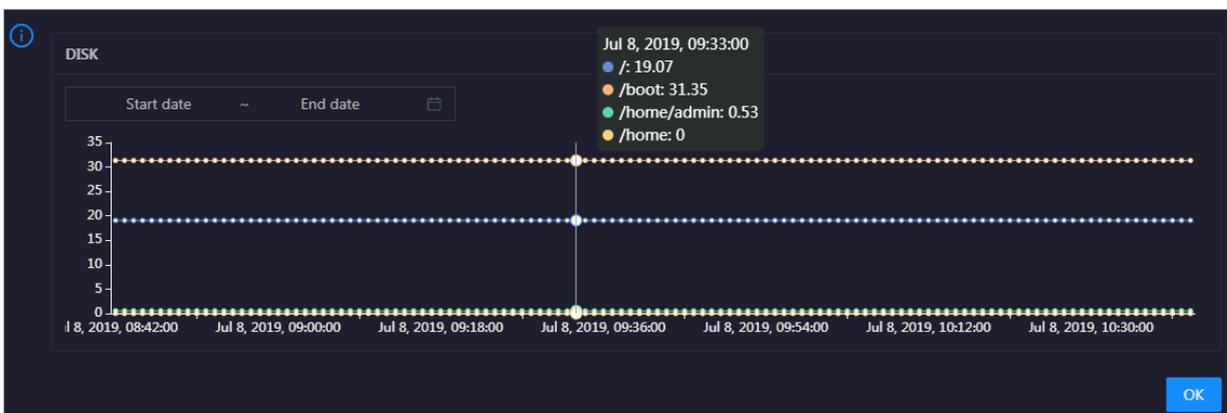


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the /, /boot, /home/admin, and /home directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



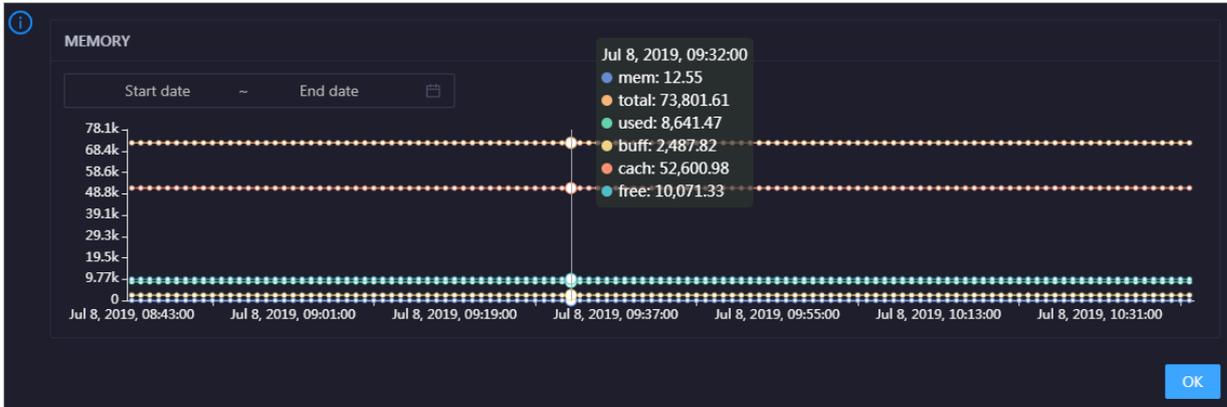
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size

of memory used by the page cache (cach), and available memory size (free) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

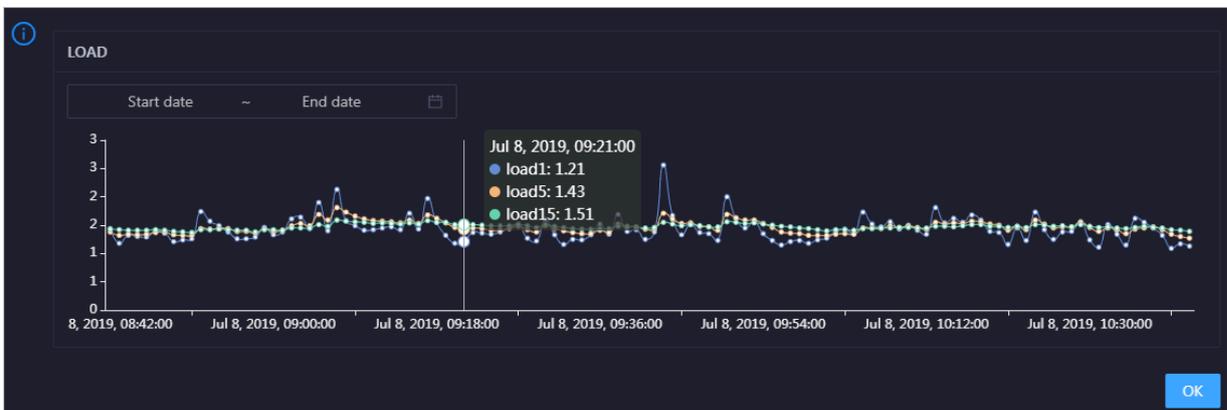


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



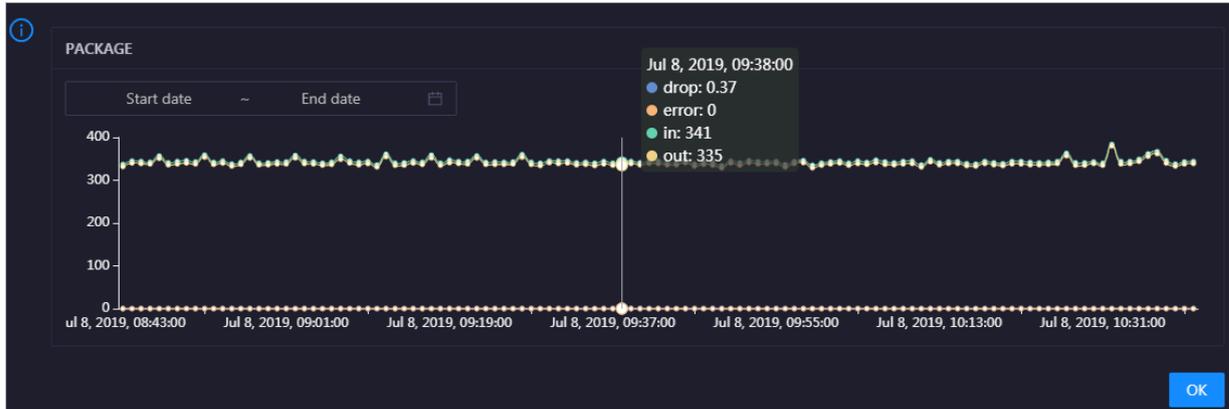
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for

the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.

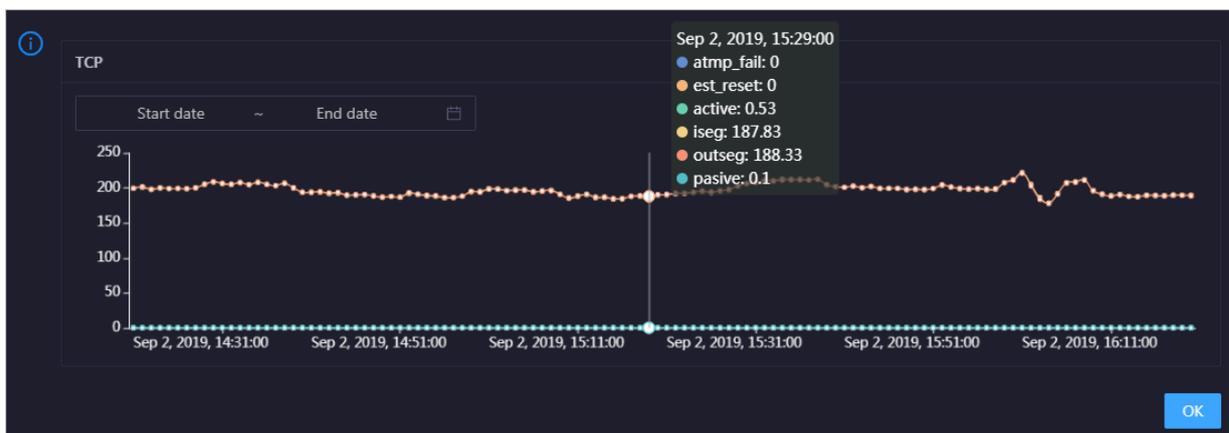


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click  in the upper-right corner of the chart to zoom in the chart.

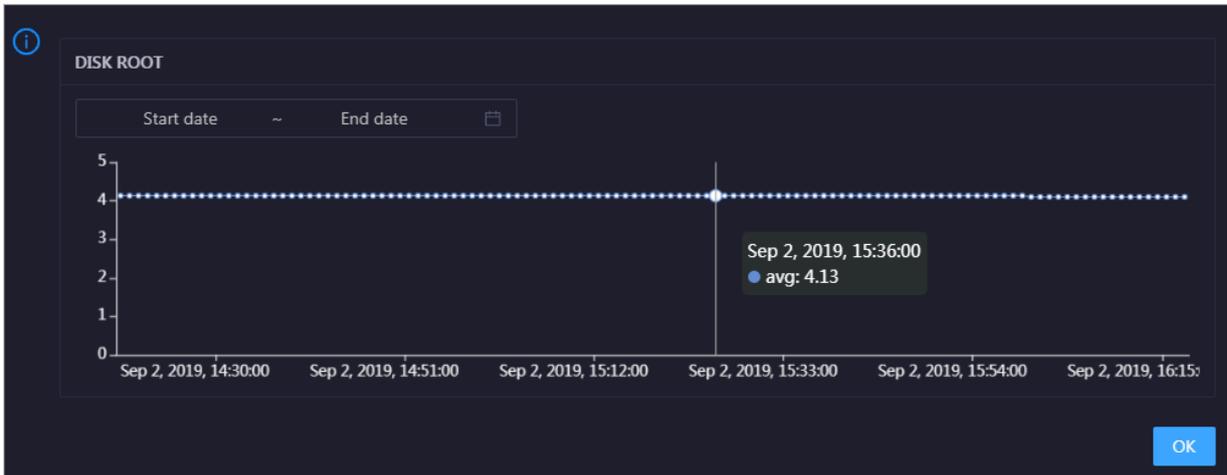


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the host over time.

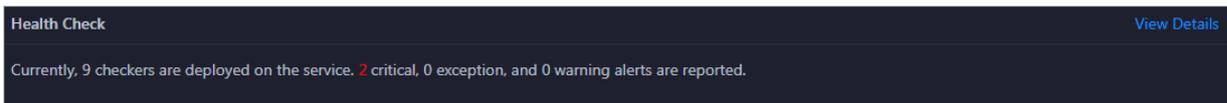
Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

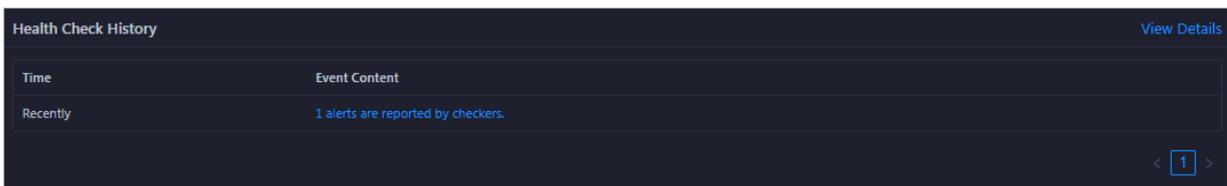
This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.



Click [View Details](#) to go to the [Host health](#) page. On this page, you can view the health check details.

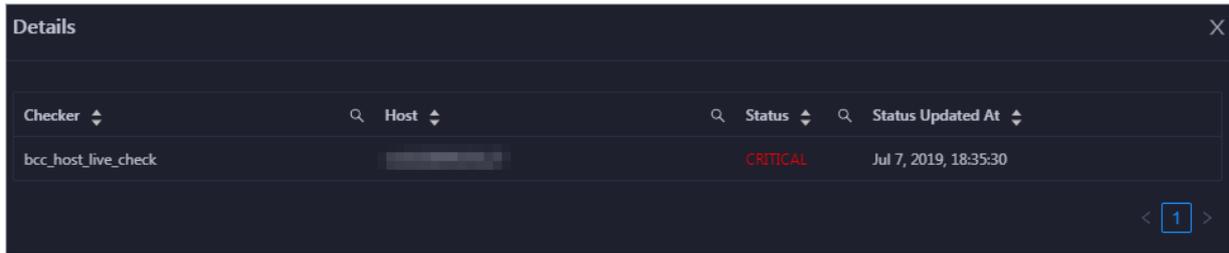
Health Check History

This section displays a record of the health checks performed on the host.



Click [View Details](#) to go to the [Host health](#) page. On this page, you can view the health check details.

You can click the event content of a check to view the exception items.



3.1.4.6.2 Host health

On the host health status page, you can view the checkers of all hosts, checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

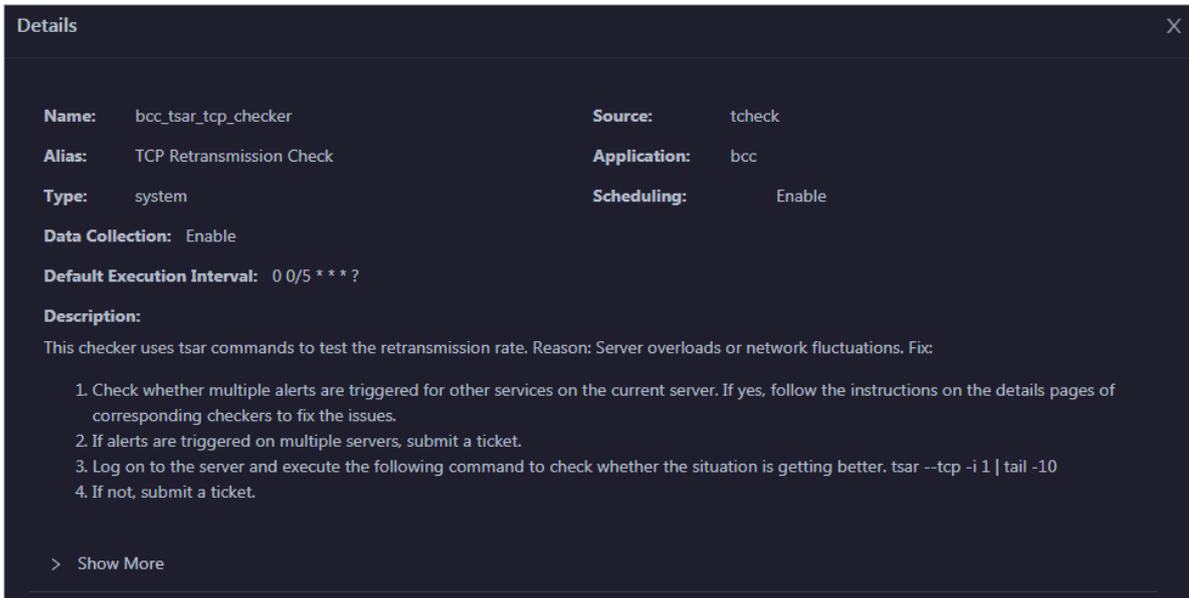
At the top of the O&M page, click Hosts. In the left-side navigation pane, select a host, and then click the Health Status tab. The Health Status page of the host appears.

Checker	Source	Critical	Warning	Exception	Actions
+ bcc_disk_usage_checker	tcheck	1	0	0	Details
+ bcc_check_ntp	tcheck	0	0	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_network_tcp_connections_checker	tcheck	0	0	0	Details
+ bcc_host_live_check	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception events are alerts. You need to pay attention to them, especially the Critical and Warning events.

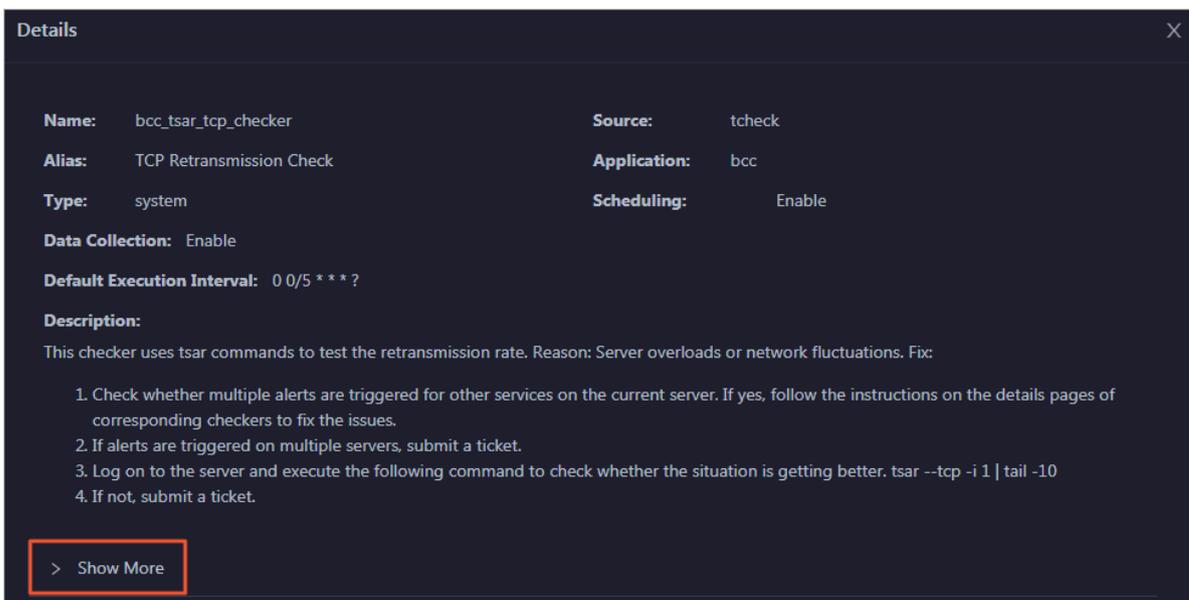
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

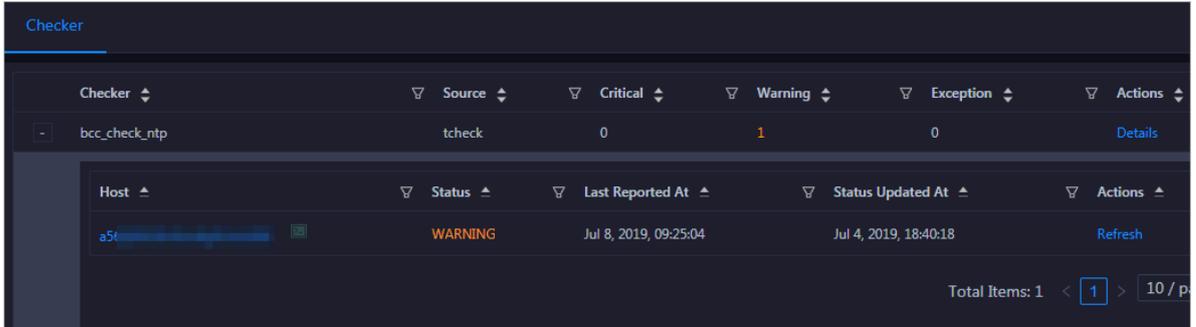


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

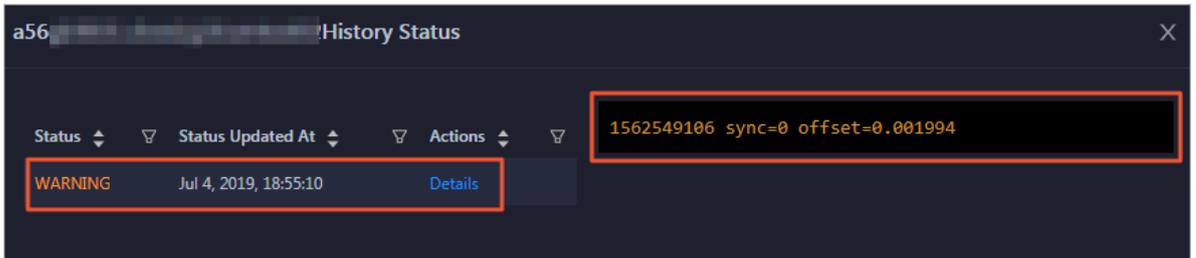
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

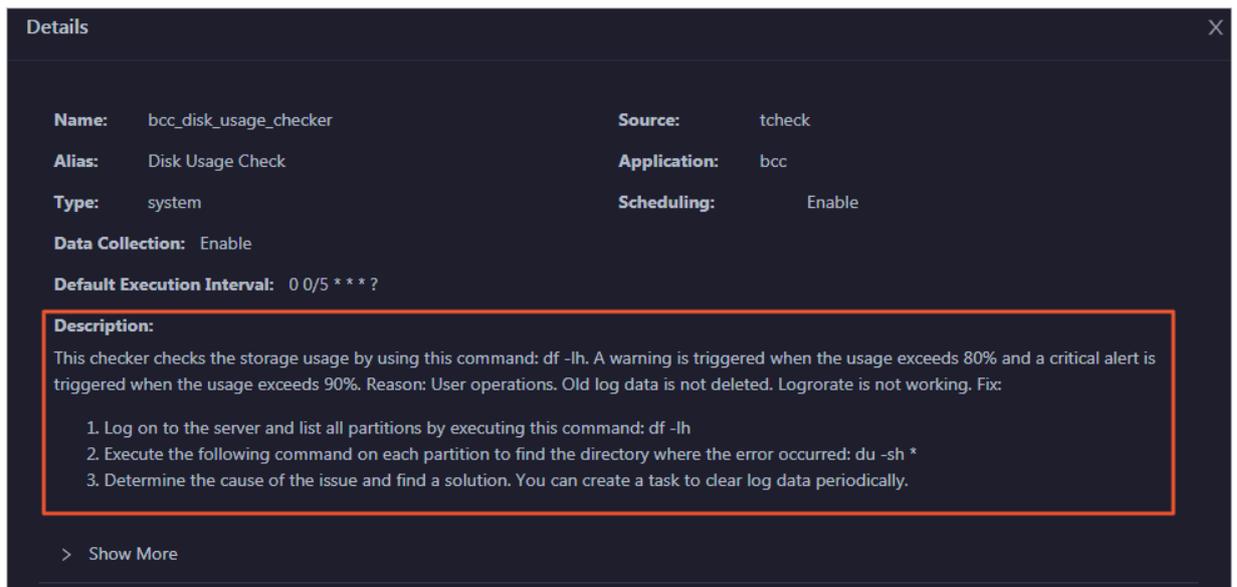


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

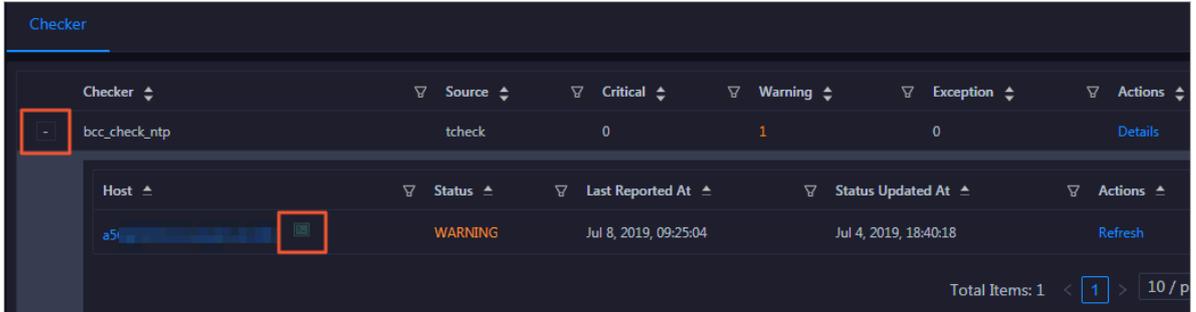
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



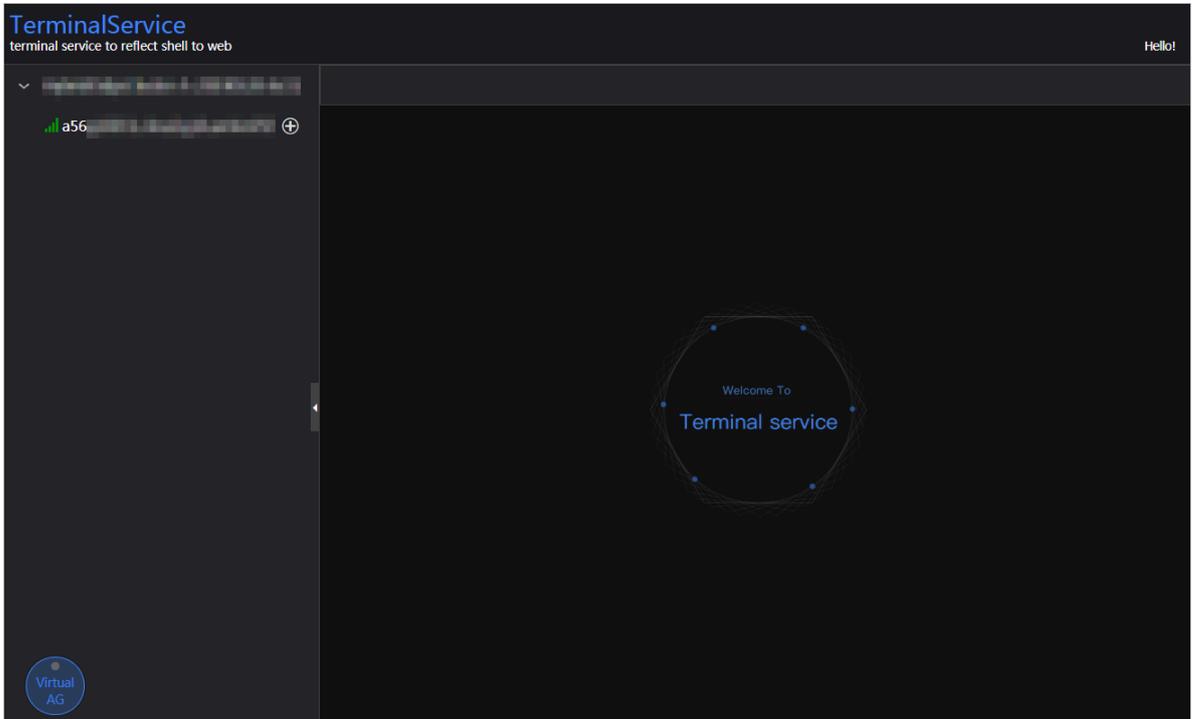
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

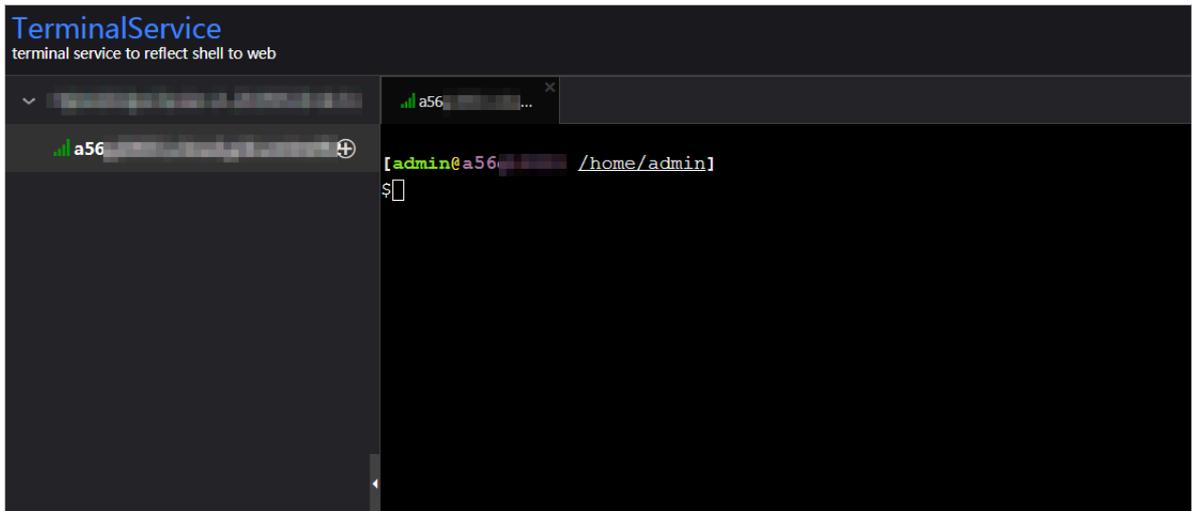
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

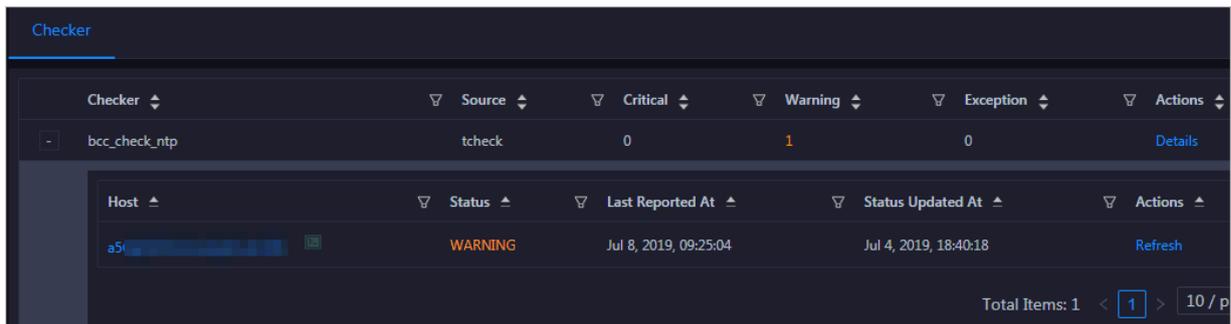


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.5 MaxCompute

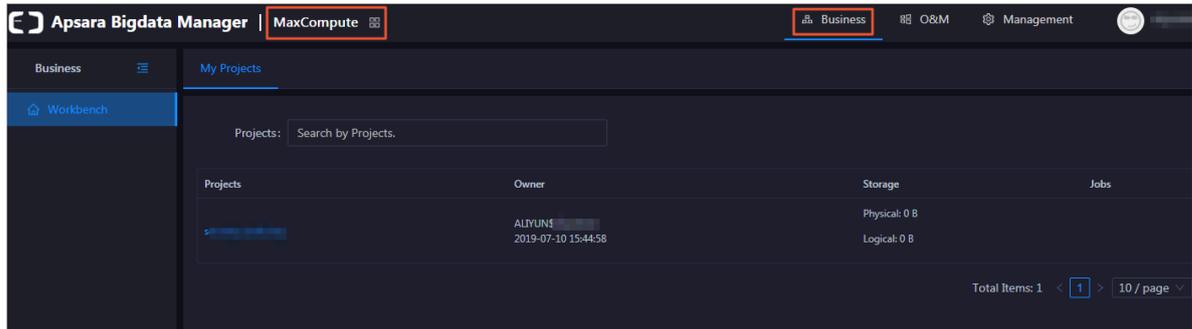
3.1.5.1 MaxCompute workbench

In the Apsara Bigdata Manager (ABM) console, you can view your MaxCompute projects and project details, including the project overview, jobs, storage, configurations, quota groups, tunnels, and resource analysis, on the MaxCompute workbench. You can also modify the current configuration of a project.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click MaxCompute.

3. On the MaxCompute page, click **Business** in the upper-right corner. The **My Projects** page under **Workbench** appears.

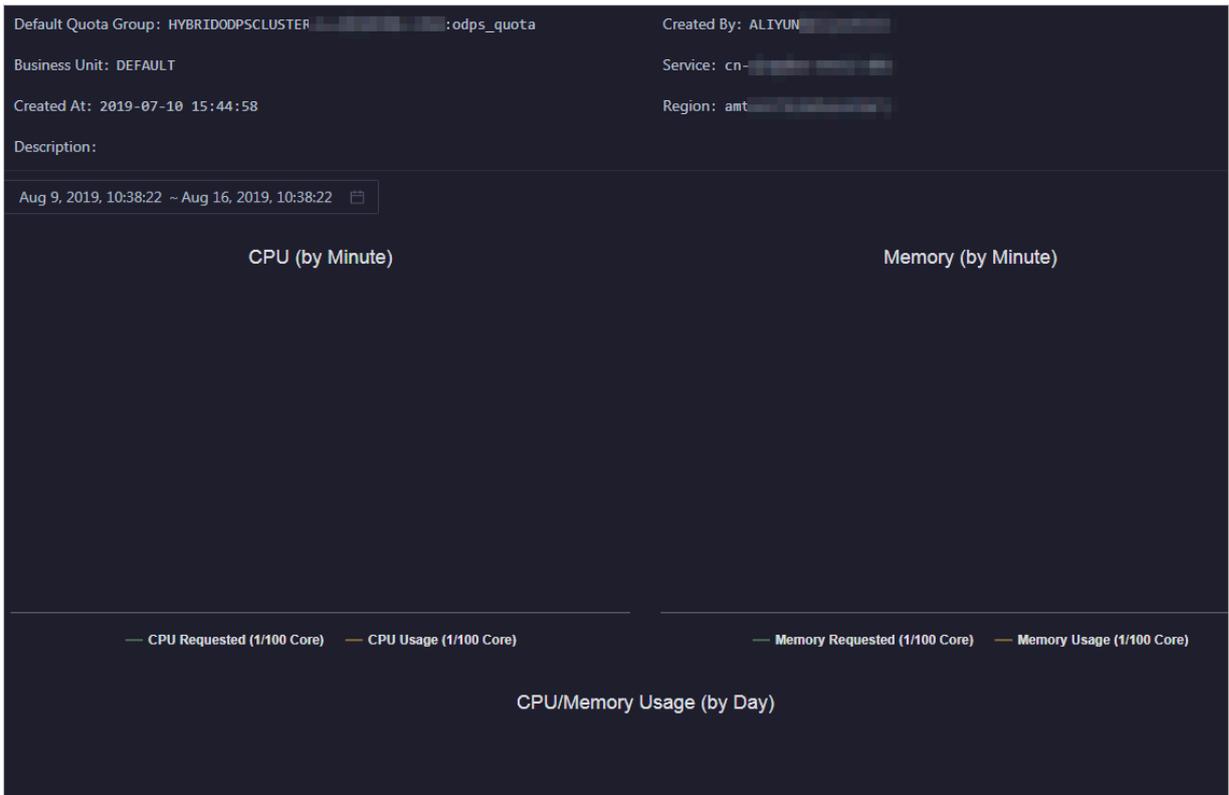


Overview

On the Overview page, you can view the following information about the selected project:

- Basic information about the project, such as the default quota group, creator, creation time, service, and region.
- Trend charts that display the trend lines of requested and used CPU and memory resources by minute over time in different colors.
- Trend chart that displays the trend lines of CPU and memory usage by day over time in different colors.

On the My Projects page, click the name of a project in the project list. The Overview page appears.



Jobs

On the Jobs page, you can view job snapshots by day over the last week. Detailed information about a job snapshot includes the job ID, project, quota group, submitter, running duration, minimum and maximum CPU usage, minimum and maximum memory usage, DataWorks node, running status, start time, priority, and type. You can also view the operational logs of a job to locate job failures.

On the My Projects page, click the name of a project in the project list and then click the Jobs tab. The Jobs page appears.

All		Running		Waiting for Resources		Initializing						
2		2		0		0						
Filter	Terminate Job	Jul 25, 2019, 16:40:39						Refresh				
JobId	Project	Quota ...	Submit...	Elapse...	CPU Us...	Memor...	DataW...	Cluster	Status	Start Tl...	Priority	Type
<input type="checkbox"/>	201907250837	odps_smoke_tr	odps_quota	ALYUN\$	18Seconds	200(200%/0.64)	2816(275%/0.2)	HYBRIDODPSC	Running	2019-07-25 16	1	CUPID
<input type="checkbox"/>	201907221435	biggraph_inter	biggraph_quot	ALYUN\$	66Hours2Minu	0(0%/0%)	0(0%/0%)	HYBRIDODPSC	Running	2019-07-22 22	1	CUPID

You can perform the following operations on jobs:

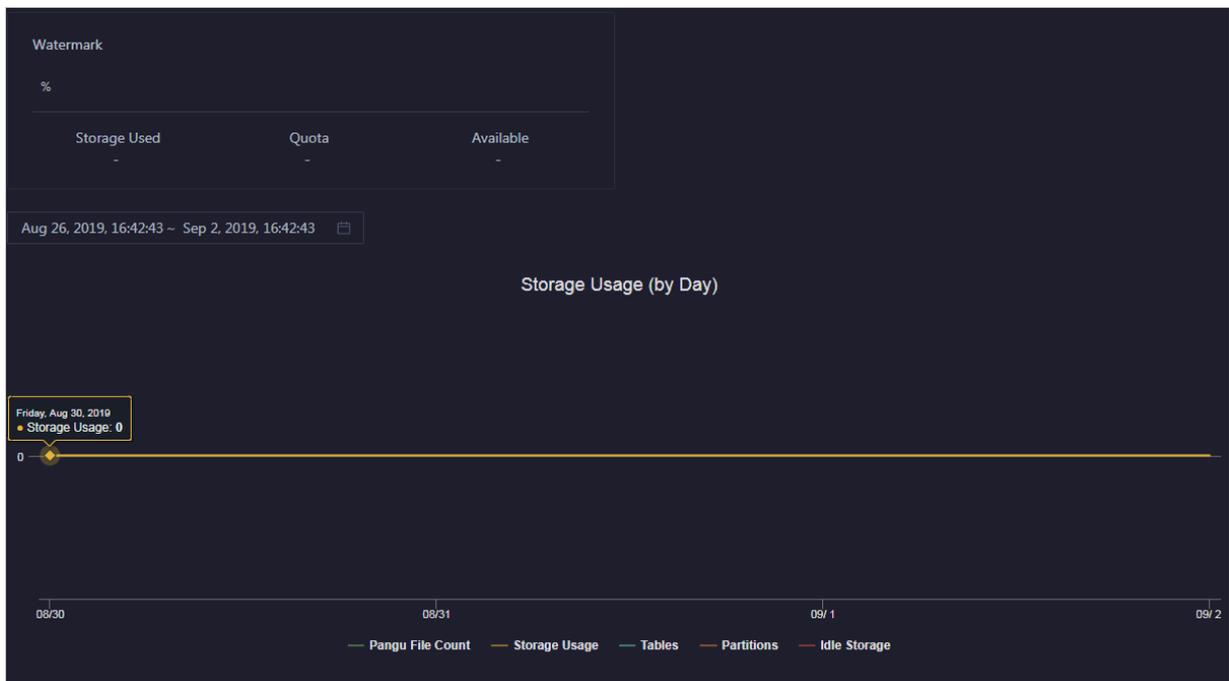
- **Customize columns or sort job snapshots based a specified column. For more information, see [Common operations](#).**

- **View operational logs of jobs or terminate jobs.** For more information, see [Job snapshots](#).

Storage

On the Storage page, you can view the storage usage, used storage space, storage quota, and available storage space. You can also view a trend chart that displays the trend lines of storage usage, the number of Apsara Distributed File System files, the number of tables, the number of partitions, and idle storage by day over time in different colors.

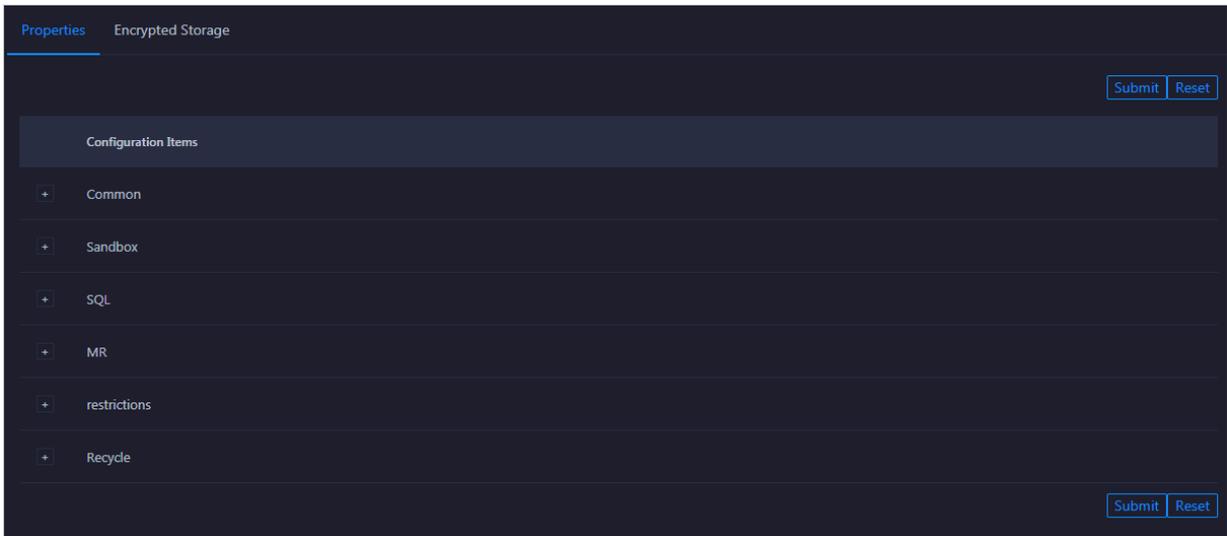
On the My Projects page, click the name of a project in the project list and then click the Storage tab. The Storage page appears.



Configuration

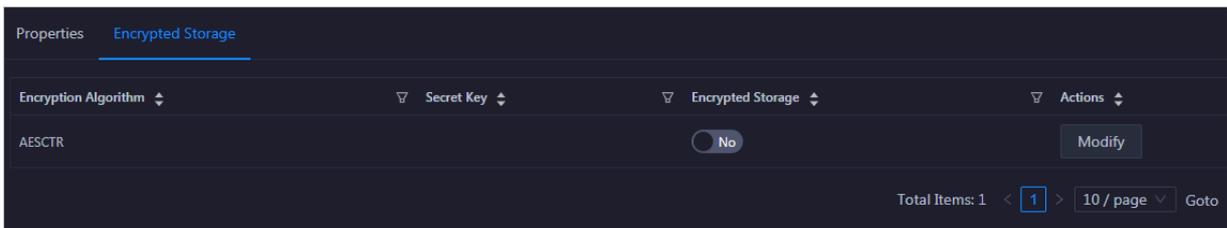
On the Configuration page, you can configure the general, sandbox, SQL, MR, access control, and resource recycling properties of the project. You can also configure encryption algorithms for the project.

On the My Projects page, click the name of a project in the project list and then click the Configuration tab. The Properties page appears.



On the Properties page, you can view and modify each configuration item. To restore all configuration items to the default settings, click Reset.

On the Encrypted Storage page, you can configure the RC4 and AESCTR encryption algorithms.



Quota groups

On the Quota Groups page, you can view the quota groups of the project and the details of each quota group.

On the My Projects page, click the name of a project in the project list and then click the Quota Groups tab. The Quota Groups page appears.

Cluster	Quota Group	Default	CPU Usage/Minimum Quota	Memory Usage/Minimum Quota	CPU Usage Percentage	Memory Usage Percentage
HYBR		Default	0 / 100	0 / 1024	0 %	0 %

To view detailed information about a quota group, click the quota group name. For more information, see [View quota group details](#).

Tunnel

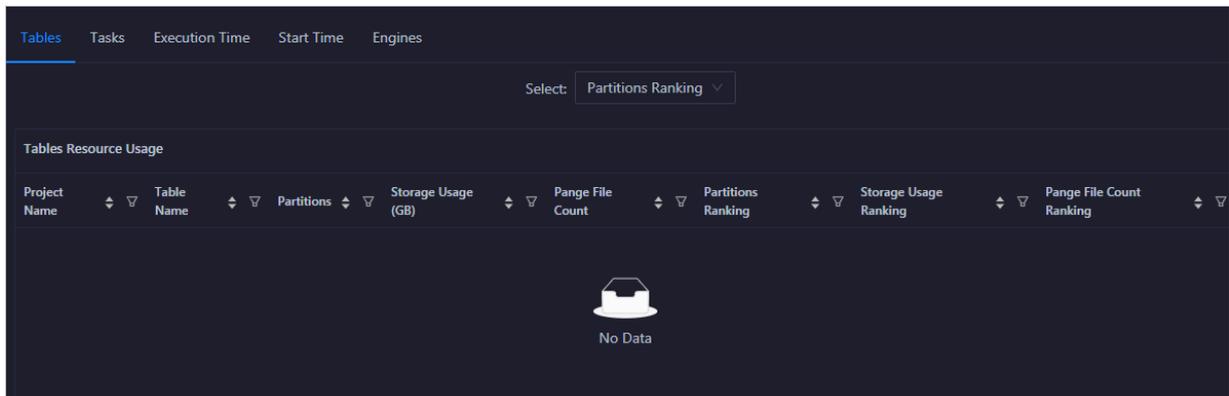
On the Tunnel page, you can view the tunnel throughput of the project in the unit of bytes per minute. The Tunnel Throughput chart displays the trend lines of inbound traffic and outbound traffic over time in different colors.

On the My Projects page, click the name of a project in the project list and then click the Tunnel tab. The Tunnel page appears.

Resource analysis

On the Resource Analysis page, you can view the resource usage for the project from different dimensions, including tables, tasks, execution time, start time, and engines. For more information, see [Resource analysis](#).

On the My Projects page, click the name of a project in the project list and then click the Resource Analysis tab. The Tables page appears.



3.1.5.2 Business O&M

3.1.5.2.1 Business O&M overview

This topic describes the features of MaxCompute business O&M and how to access the MaxCompute business O&M page.

Modules

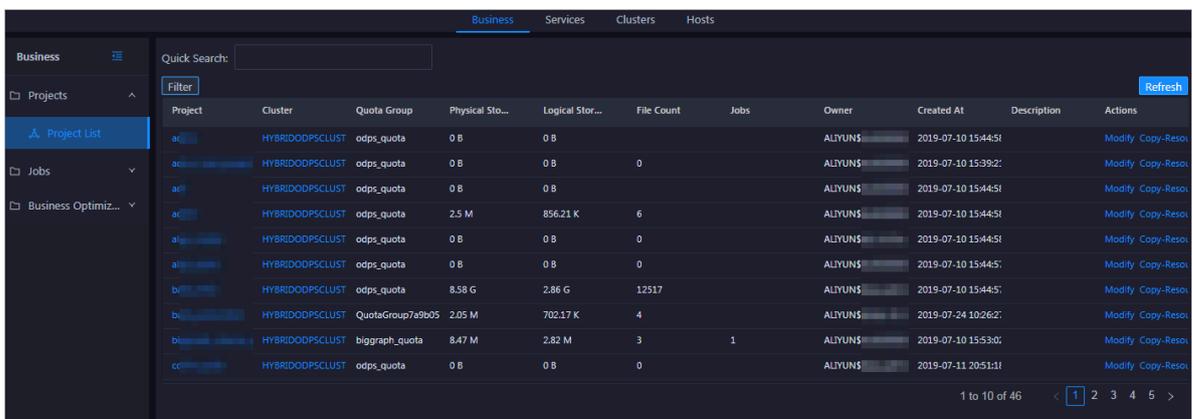
- **Project management:**
 - **Project list:** displays all projects and project details in the MaxCompute cluster. You can filter, query, and sort projects. You can also modify the quota group of a project. If zone-disaster recovery is enabled, you can set resource

replication parameters and determine whether to enable resource replication for a project.

- **Disaster recovery:** allows you to view the cluster status when zone-disaster recovery is enabled for MaxCompute. You can enable switchover between the primary and standby clusters. You can also determine whether to run scheduled tasks to synchronize resources between the primary and standby clusters.
- **Job management:** displays information about jobs deployed on a MaxCompute cluster. You can filter and search for these jobs. You can also view the operation logs, terminate a running job, and collect job logs.
- **Business optimization:**
 - **File merging:** allows you to create file merge tasks for clusters and projects. You can also filter merge tasks and view records of the tasks.
 - **File archiving:** allows you to create file archiving tasks for clusters and projects. You can also filter archiving tasks and view records of the tasks.
 - **Resource analysis:** allows you to view the resource usage for the cluster from different dimensions.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page, click O&M in the upper-right corner, and then click Business. The Project List page under Projects appears.



Project	Cluster	Quota Group	Physical Sto...	Logical Stor...	File Count	Jobs	Owner	Created At	Description	Actions
ac...	HYBRIDODPSCLUST	odps_quota	0 B	0 B			ALYUN\$...	2019-07-10 15:44:51		Modify Copy-Resol
ad...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$...	2019-07-10 15:39:21		Modify Copy-Resol
ae...	HYBRIDODPSCLUST	odps_quota	0 B	0 B			ALYUN\$...	2019-07-10 15:44:51		Modify Copy-Resol
af...	HYBRIDODPSCLUST	odps_quota	2.5 M	856.21 K	6		ALYUN\$...	2019-07-10 15:44:51		Modify Copy-Resol
ag...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$...	2019-07-10 15:44:51		Modify Copy-Resol
ah...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$...	2019-07-10 15:44:51		Modify Copy-Resol
ai...	HYBRIDODPSCLUST	odps_quota	8.58 G	2.86 G	12517		ALYUN\$...	2019-07-10 15:44:51		Modify Copy-Resol
aj...	HYBRIDODPSCLUST	QuotaGroup7a9b05	2.05 M	702.17 K	4		ALYUN\$...	2019-07-24 10:26:21		Modify Copy-Resol
ak...	HYBRIDODPSCLUST	biggraph_quota	8.47 M	2.82 M	3	1	ALYUN\$...	2019-07-10 15:53:01		Modify Copy-Resol
al...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$...	2019-07-11 20:51:11		Modify Copy-Resol

3.1.5.2.2 Project management

3.1.5.2.2.1 Project list

The Project List page displays all projects and project details in the MaxCompute cluster. You can filter, query, and sort projects. You can also modify the quota group of a project. If zone-disaster recovery is enabled, you can set resource replication parameters and determine whether to enable resource replication for a project.

Project List page

On the Business page, choose Projects > Project List in the left-side navigation pane to view projects in the cluster.

Project	Cluster	Quota Group	Physical Sto...	Logical Stor...	File Count	Jobs	Owner	Created At	Description	Actions
ac...	HYBRIDODPSCLUST	odps_quota	0 B	0 B			ALYUN\$	2019-07-10 15:44:51		Modify Copy-Reso...
ad...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$	2019-07-10 15:39:21		Modify Copy-Reso...
ae...	HYBRIDODPSCLUST	odps_quota	0 B	0 B			ALYUN\$	2019-07-10 15:44:51		Modify Copy-Reso...
af...	HYBRIDODPSCLUST	odps_quota	2.5 M	856.21 K	6		ALYUN\$	2019-07-10 15:44:51		Modify Copy-Reso...
ag...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$	2019-07-10 15:44:51		Modify Copy-Reso...
ah...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$	2019-07-10 15:44:51		Modify Copy-Reso...
ai...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$	2019-07-10 15:44:51		Modify Copy-Reso...
aj...	HYBRIDODPSCLUST	odps_quota	8.58 G	2.86 G	12517		ALYUN\$	2019-07-10 15:44:51		Modify Copy-Reso...
ak...	HYBRIDODPSCLUST	QuotaGroup7a9b05	2.05 M	702.17 K	4		ALYUN\$	2019-07-24 10:26:21		Modify Copy-Reso...
al...	HYBRIDODPSCLUST	biggraph_quota	8.47 M	2.82 M	3	1	ALYUN\$	2019-07-10 15:53:01		Modify Copy-Reso...
am...	HYBRIDODPSCLUST	odps_quota	0 B	0 B	0		ALYUN\$	2019-07-11 20:51:11		Modify Copy-Reso...

On the Project List page, you can view detailed information about all projects in the cluster. For example, you can view the name, cluster, storage, file quantity, running job quantity, owner, creation time, quota group, and description of a project.

Facilitate information retrieval

You can filter and query projects. You can also customize columns or sort projects based a specified column. For more information, see [Common operations](#).

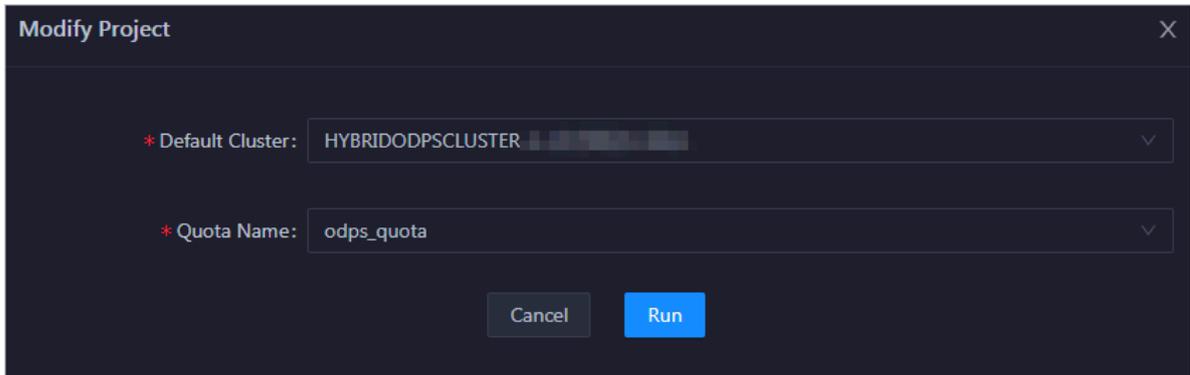
View project details

On the Project List page, you can click the name of a project to view its detailed information, including the project overview, jobs, storage, configurations, quota groups, tunnels, and resource analysis. For more information, see [MaxCompute workbench](#).

Modify a project

You can modify the quota group and default cluster of a project.

1. On the Project List page, find the project to be modified and click **Modify** in the Actions column. In the Modify Project dialog box that appears, set related parameters.



The screenshot shows a dark-themed dialog box titled "Modify Project". It features two dropdown menus. The first is labeled "* Default Cluster:" and has "HYBRIDODPSCLUSTER" selected. The second is labeled "* Quota Name:" and has "odps_quota" selected. At the bottom of the dialog, there are two buttons: "Cancel" and "Run".

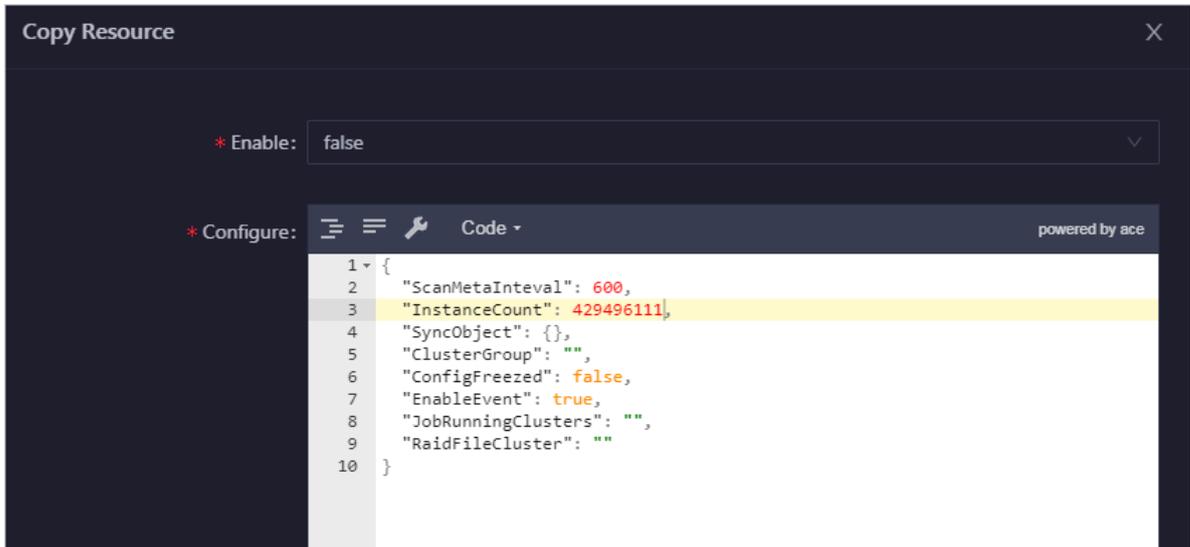
The parameters are described as follows:

- **Default Cluster:** the default cluster of the project. If the project belongs to multiple clusters, you can select a cluster from the drop-down list to serve as the default cluster.
 - **Quota Name:** the quota group to which the project belongs. To change the quota group, select the specified quota group from the drop-down list.
2. Click **Run**. A message appears, indicating that the action has been submitted.

Configure resource replication

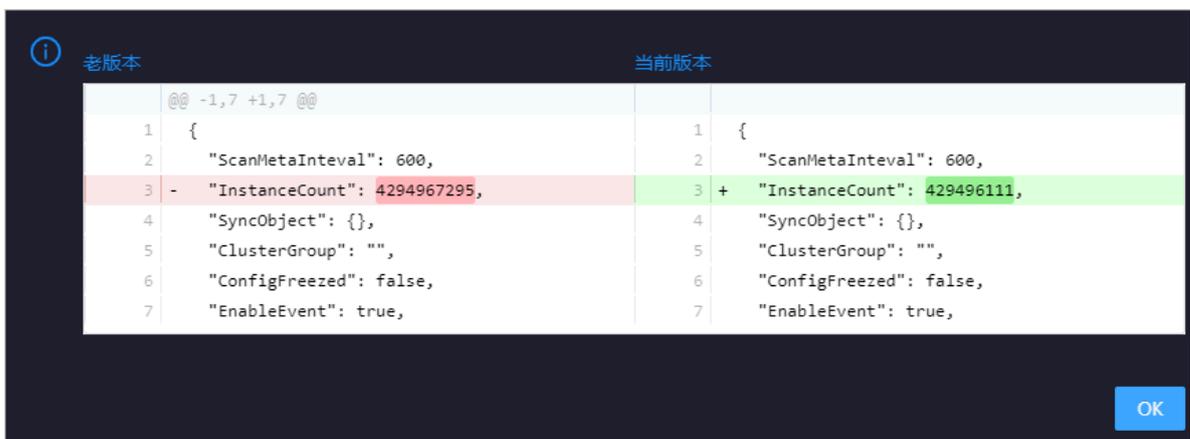
The resource replication feature can be configured only in zone-disaster recovery scenarios. In other scenarios, you can only view the settings. In zone-disaster recovery scenarios, you can determine whether to enable the resource replication feature for a project in the primary cluster. If the resource replication feature is enabled for a project, you can set data synchronization rules for the project to regularly synchronize data such as data tables to the standby cluster.

1. On the Project List page, find a specified project and click Copy-Resource in the Actions column. In the Copy Resource dialog box that appears, set related parameters.



The parameters are described as follows:

- **Enable:** specifies whether to enable the resource replication feature. A value of true indicates that the resource replication feature is enabled. A value of false indicates that the resource replication feature is disabled.
 - **Configure:** the data synchronization rules of a project. Generally, the default settings are used. If you need to modify the settings, consult second-line O&M engineers.
2. After modifying the code in Configure, click Compare Versions to view the highlighted differences between the code of the current version and that of the earlier version.



3. Click Run. A message appears, indicating that the action has been submitted.

3.1.5.2.2.2 Disaster recovery

When the primary MaxCompute cluster fails, you can quickly switch services from the primary cluster to the standby cluster in the Apsara Bigdata Manager (ABM) console to restore services. This topic describes the elements on the Disaster Recovery page and the prerequisites and procedure for switchover. Only zone-disaster recovery is supported.

Disaster Recovery page



Note:

The disaster recovery feature is available only when zone-disaster recovery is enabled.

On the Business page, choose Projects > Disaster Recovery in the left-side navigation pane to view information about the primary and standby MaxCompute clusters.

The Disaster Recovery page consists of the following elements:

- **Primary and Standby sections:** The Primary section displays the name of the primary cluster and the number of projects with the primary cluster as the default cluster. The Standby section displays the name of the standby cluster and the number of projects with the standby cluster as the default cluster.



Note:

The total number of projects in the primary cluster is the same as that in the standby cluster.

- **Resource Synchronization Status:** specifies whether to run the scheduled task to synchronize resources between the primary and standby MaxCompute clusters. When you turn on this switch, resources are scheduled between the primary and standby clusters every 30 minutes.
- **View Resource Synchronization History:** allows you to view the execution history of the scheduled resource synchronization task.

Prerequisites for switchover

- You have built a zone-disaster recovery environment.
- You have obtained an ABM account with MaxCompute O&M permissions and can log on to the ABM console.

- **The VIP address of the current ABM cluster has been switched to the standby ABM cluster. For more information, see [Switch the VIP address the current ABM cluster to the standby ABM cluster](#).**
- **You have disabled the scheduled task for synchronizing resources between the primary and standby MaxCompute clusters. For more information, see [Enable or disable the resource synchronization between primary and standby MaxCompute clusters](#).**
- **The Business Continuity Management Center (BCMC) switchover of MaxCompute has been completed. The services on which MaxCompute depends, including AAS, Table Store, and MiniRDS, are running properly.**

Enable or disable the resource synchronization between primary and standby MaxCompute clusters

When the resource synchronization feature is enabled, the scheduled resource synchronization task is run to synchronize resources, such as a compiled JAR package, between the primary and standby MaxCompute clusters every 30 minutes. You need to keep the scheduled resource synchronization task disabled until the switchover between the primary and standby clusters is completed.

- 1. On the Business page, choose Projects > Disaster Recovery in the left-side navigation pane to view information about the primary and standby MaxCompute clusters.**

If the Resource Synchronization Status switch is turned on, the scheduled resource synchronization task is enabled. If the switch is turned off, the task is disabled.

- 2. Turn on or off the Resource Synchronization Status switch to enable or disable the scheduled resource synchronization task between the primary and standby clusters.**

Switch the VIP address the current ABM cluster to the standby ABM cluster

Before the switchover between the primary and standby MaxCompute clusters, you can execute the Change Bcc Dns-Vip Relation For Disaster Recovery scheme in ABM to replace the VIP address of the standby ABM cluster with that of the current ABM cluster.

- 1. [Log on to the ABM console](#).**
- 2. Click  in the upper-left corner, and then click MaxCompute.**

3. On the MaxCompute page that appears, click Management in the upper-right corner. The Management appears.
4. Click Jobs in the left-side navigation pane, and click Job Management on the right side to go to the Schemes page.
5. In the scheme list, find the Change Bcc Dns-Vip Relation For Disaster Recovery scheme, and click Run in the Actions column. On the page that appears, set the following two parameters in Target Group:

Set the NowBccApiOneIp parameter to the IP address of any Docker container in a path of the current ABM cluster, for example, bcc > bcc-api > controller# > #Docker#xx.xx.xx.xx. Set the NewBccApiOneIp parameter to the IP address of any Docker container in the same path of the standby ABM cluster.
6. Click Run in the upper-right corner and confirm the risks of running the job.
7. Click Confirm. The job running page appears.
8. Click Start at the top of the page.

Start switchover

After all the prerequisites for switchover are met, you can start MaxCompute switchover.

1. On the Business page, choose Projects > Disaster Recovery in the left-side navigation pane to view information about the primary and standby MaxCompute clusters.
2. Click Start Switchover in the upper-right corner. A dialog box appears, asking you to confirm whether the VIP address of the standby ABM cluster has been replaced with that of the current ABM cluster.

If the VIP address of the standby ABM cluster has not been replaced with that of the current ABM cluster, click No and then replace the VIP address of the standby ABM cluster with that of the current ABM cluster. For more information, see

[Switch the VIP address the current ABM cluster to the standby ABM cluster.](#)

3. Click Yes. The Stop Resource Replication page appears.



Note:

In this step, the scheduled resource synchronization task is automatically disabled.

4. After resource replication is disabled, click **Next Step**. The **Switch Control Cluster** page appears.

In this step, the services are automatically switched from the primary cluster to the standby cluster, which takes about 30 seconds. You can determine whether the switchover is successful based on the values of the **Current Primary Cluster** and **Current Standby Cluster** parameters on the page. When the primary and standby clusters are switched, the switchover is complete.

After the switchover, you need to perform the following operations:

- a. Click **Restart Standby Cluster**. It takes about 20 seconds to restart the standby cluster. When the standby cluster is restarted, the value of the **MaxCompute Cluster Status** parameter changes from **Abnormal** to **Normal**.
 - b. Click **Restart Frontend Server**. It takes about 20 seconds to restart the front-end server. When the front-end server is restarted, a success message appears.
 - c. Click **Test adminTask** to check whether the **MaxCompute** service is normal. If the test is passed, the clusters are switched. The **Next Step** button becomes operable, and the **Switching...** message disappears.
5. Click **Next Step**. The **Switch Computing Cluster** page appears.

In this step, the default computing cluster of the projects in the primary cluster is changed to the standby cluster, and that of the projects in the standby cluster is changed to the primary cluster. Each project has a switchover progress bar. If the progress bar of a project is highlighted, the switchover is complete.



Note:

If the computing cluster of a project fails to be switched, you can contact O&M engineers to locate the cause of the exception. If the project can be fixed, fix it and click **Retry** to continue the switchover. If the project is damaged or does not need to change the computing cluster, you can click **Next Step** after confirming that other projects have been switched.

6. Click **Next Step**. The **Switch Replication Service to Standby Service** page appears.

7. After the switchover is completed, click **Next Step**. The **Collect Statistics about Unsynchronized Data** page appears.

This step takes some time, depending on the data volume. Wait until the step is completed. After the collection is completed, the system lists all projects with unsynchronized data. You can check the data that has not been synchronized.

You must select the projects with unsynchronized data, and click **Download Unsynchronized Data of Selected Projects** to download the data to a local device so that you can manually fill in the missing data later based on the statistics. Only after the unsynchronized data is downloaded does the **Next Step** button become operable.



Note:

If the unsynchronized data is abnormal, you can click **Recollect Unsynchronized Data**.

8. Click **Next Step**. The **Repair Metadata** page appears.

In this step, the data in the primary and standby clusters becomes the same. Select all projects, click **Repair Metadata of Selected Projects**, and then wait for results.

- If some projects fail to be fixed, click **Download Last Execution Log** and send the logs to O&M engineers to analyze the cause of the exception. After the exception is resolved, you can fix the projects again.
- If you do not need to fix all projects, click **Next Step** after the necessary projects are fixed.

9. After the metadata is fixed, click **Next Step**. The **Manually Fill in Missing Data** page appears.

In this step, you need to log on to the DataWorks console, and manually fill in the missing data according to the unsynchronized data downloaded in the **Collect Statistics about Unsynchronized Data** step. After filling in the missing data, select all projects and click **Confirm Data Repair Complete**. Then, the **Next Step** button becomes operable.

10. Click **Next Step**. The **Repair Unsynchronized Resources** page appears.

In this step, it takes some time to count the projects to be fixed, depending on the data volume. Wait until the results are displayed. If some projects in the standby

cluster are inconsistent with those in the primary cluster, you need to fill in the missing data manually. Otherwise, proceed to the next step.

11. After the unsynchronized resources are fixed, click **Complete** and **Next**. The **Enable Resource Replication** page appears.

In this step, the scheduled resource synchronization task is automatically started .

12. After enabling resource replication, click **Next Step**. The **Complete Wizard** page appears.

13. Click **Back** in the upper-left corner. The primary and standby clusters have been switched.

3.1.5.2.3 Job management

3.1.5.2.3.1 Job snapshots

The **Job Snapshots** page allows you to manage the tasks created in MaxCompute and the merge tasks created in Apsara Bigdata Manager (ABM). You can also view the job details by using Logview, terminate a job, and collect job logs on this page.

View job snapshots

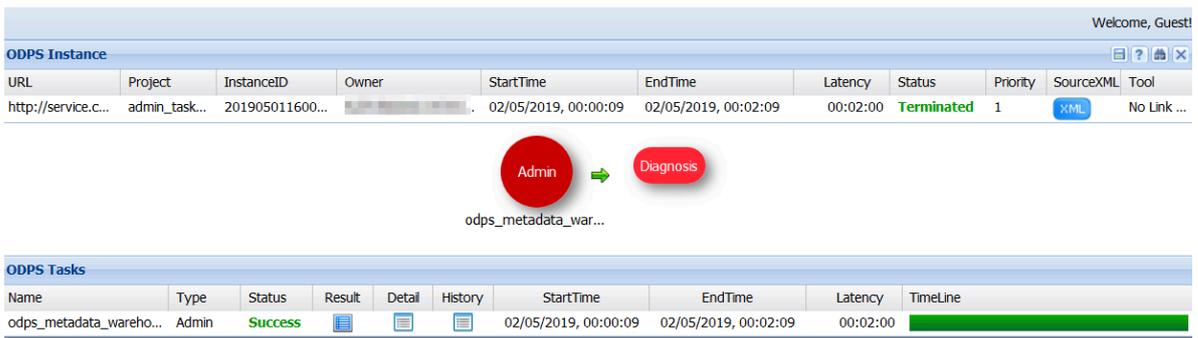
You can view job snapshots by day over the last week. Detailed information about a job snapshot includes the job ID, project, quota group, submitter, running duration , minimum and maximum CPU usage, minimum and maximum memory usage, DataWorks node, running status, start time, priority, and type. You can also view the operational logs of a job to locate job failures.

The **Job Snapshots** page supports multiple operations to facilitate information retrieval. For example, you can filter and sort job snapshots. For more information, see [Common operations](#).

1. On the **Business** page, choose **Jobs > Job Snapshots** in the left-side navigation pane. The **Job Snapshots** page appears.

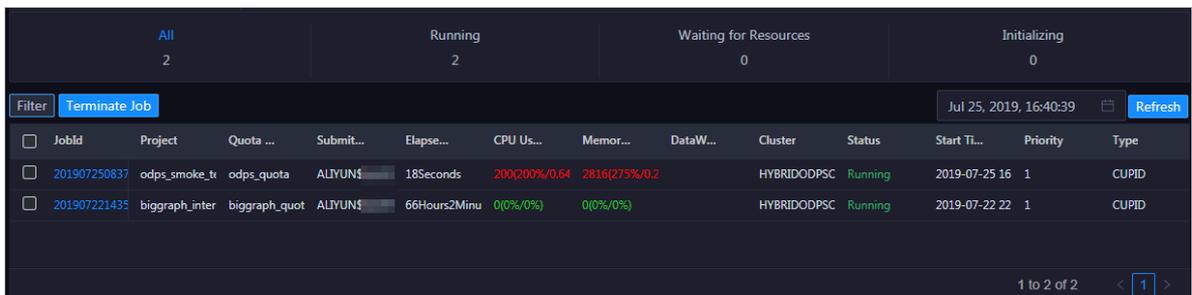
JobId	Project	Quota ...	Submit...	Elapse...	CPU Us...	Memor...	DataW...	Cluster	Status	Start Ti...	Priority	Type
201907250837	odps_smoke_tr	odps_quota	ALIYUN\$	18Seconds	200(200%/0.64)	2816(275%/0.2)		HYBRIDODPSC	Running	2019-07-25 16	1	CUPID
201907221435	biggraph_inter	biggraph_quot	ALIYUN\$	66Hours2Minu	0(0%/0%)	0(0%/0%)		HYBRIDODPSC	Running	2019-07-22 22	1	CUPID

2. In the upper-right corner of the job snapshot list, select the date and time to view job snapshots by day over the last week.
3. Click All, Running, Waiting for Resources or Initializing to view job snapshots in the corresponding status on the specified date.
4. Click the job ID of a job snapshot, and then click DetailLogview. A dialog box appears, containing a link to Logview.
5. Click click here to access the Logview page and view detailed information about the job.

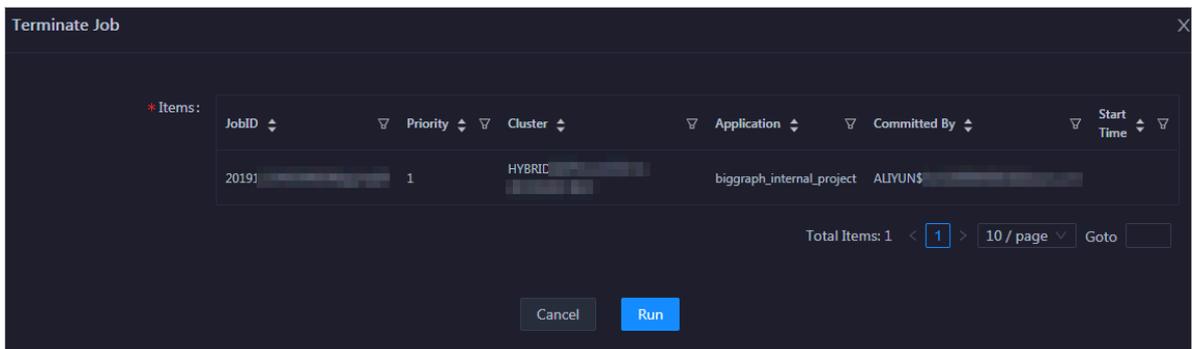


Terminate a job

1. On the Business page, choose Jobs > Job Snapshots in the left-side navigation pane. The Job Snapshots page appears.



2. Select one or more jobs, and then click Terminate Job. In the dialog box that appears, view the information about the job or jobs to be terminated.

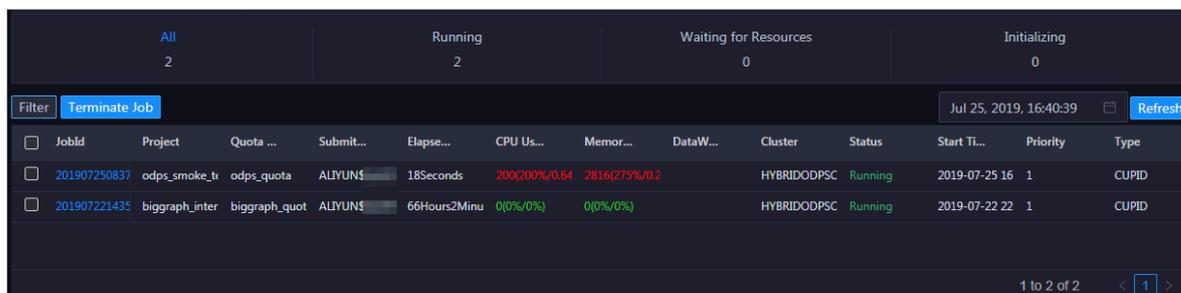


3. Click Run. A message appears, indicating the running result.

Collect job logs

When an exception occurs during job running, you can locate and analyze the issue by collecting job logs.

1. On the Business page, choose Jobs > Job Snapshots in the left-side navigation pane. The Job Snapshots page appears.



2. In the upper-left corner of the Job Snapshots, choose Actions > Collect Job Logs.
3. In the Collect Job Logs dialog box that appears, set the parameters.

The following table describes the required parameters.

Parameter	Description
Target Service	The target service from which you want to collect job logs. Select a target service from the drop-down list.
intanceid	(Optional) The ID of the job instance.
requestid	(Optional) The ID of the request returned when the job execution fails. If the value you specify is not a request ID, job logs that contain the corresponding value will be collected.
Time Period	The time range in which you want to collect job logs.
Time Interval	(Optional) The time interval for collecting job logs. Unit: hours.
Degree of Concurrency	The maximum number of nodes from which you can collect job logs at the same time.

4. Click Run to start job log collection.

5. View the execution status and progress of job log collection.

In the upper-left corner of the Job Snapshots page, choose Actions, and then click  next to Collect Job Logs. The execution status and history of job log collection are displayed.

In the Current Status column, RUNNING indicates that the execution is in progress, FAILED indicates that the execution fails, and SUCCESS indicates that the execution is successful. You can click Details in the Details column of a task in the Running state to view the execution progress.

6. View the path for storing the job logs.

You can click Details in the Details column of a collection task in the Success state to view the execution details. In the Steps section, click the name of the node to show detailed information, and then click View Details in the Actions column to view the path for storing the job logs.

3.1.5.2.4 Business optimization

3.1.5.2.4.1 File merging

Excessive small files in a MaxCompute cluster occupy a lot of memory resources. Apsara Bigdata Manager (ABM) allows you to merge small files in clusters and projects to release memory resources occupied by excessive small files.

Create a merge task for a cluster

When excessive small files exist in most projects of a MaxCompute cluster, you can create a merge task to merge the small files in the cluster in a unified manner.

1. On the Business page, choose Business Optimization > File Merging in the left-side navigation pane. The Merge Tasks tab appears.
2. In the Merge Tasks for Clusters section, click Create Merge Task, and set relevant parameters in the dialog box that appears.

The following table describes the required parameters.

Parameter	Description
Cluster	The cluster in which you want to run the merge task. Select a cluster from the drop-down list.
Start Time	The start time of the merge task.
End Time	The end time of the merge task.

Parameter	Description
Bandwidth Limit	<p>Specifies whether to support the concurrency of the merge tasks for the cluster.</p> <ul style="list-style-type: none"> • Yes: indicates that merge tasks cannot be run simultaneously. • No: indicates that merge tasks can be run simultaneously.
Maximum Concurrent Tasks	The maximum number of merge tasks that can be run at the same time in the selected cluster. This parameter is valid only when Bandwidth Limit is set to No.
Enabled	Specifies whether the merge task is enabled.
Merge Parameters	<p>The parameter configuration for the merge task. You can use the following default configuration:</p> <pre> { "odps.idata.useragent": "SRE Merge", "odps.merge.cpu.quota": "75", "odps.merge.quickmerge.flag": "true", "odps.merge.cross.paths": "true", "odps.merge.smallfile.filesize.threshold": "4096", "odps.merge.maxmerged.filesize.threshold": "4096", "odps.merge.max.filenumber.per.instance": "10000", "odps.merge.max.filenumber.per.job": "10000000", "odps.merge.maintain.order.flag": "true", "odps.merge.failure.handling": "any" } </pre>
Maximum Running Jobs	The maximum number of jobs that can be run at the same time in the selected cluster. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the selected cluster, not only the merge tasks.

3. Click **Compare Versions** to view the differences between the modified values and the original parameter values.

4. Click **Run**. A message appears, indicating that the action has been submitted.

After the merge task is created, it appears in the list of merge tasks for clusters.

Create a merge task for a project

When excessive small files exist in only a few projects of a MaxCompute cluster, you can create a merge task to merge the small files in each project.

1. On the **Business** page, choose **Business Optimization > File Merging** in the left-side navigation pane. The **Merge Tasks** tab appears.

2. In the Merge Tasks for Projects section, click Create Merge Task, and set relevant parameters in the dialog box that appears.

The following table describes the required parameters.

Parameter	Description
Region	The region where the cluster of the selected project resides. Select a region from the drop-down list.
Project Name	The name of the project in which you want to run the merge task. Select a project from the drop-down list.
Start Time	The start time of the merge task.
Priority	The priority of the merge task. A smaller value indicates a higher priority.
End Time	The end time of the merge task.
Enabled	Specifies whether the merge task is enabled.
Bandwidth Limit	Specifies whether to support the concurrency of the merge tasks for the project. <ul style="list-style-type: none"> • Yes: indicates that merge tasks cannot be run simultaneously. • No: indicates that merge tasks can be run simultaneously.
Maximum Concurrent Tasks	The maximum number of merge tasks that can be run at the same time in the cluster of the selected project. This parameter is valid only when Bandwidth Limit is set to No.
Maximum Running Jobs	The maximum number of jobs that can be run at the same time in the cluster of the selected project. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the cluster of the selected project, not only the merge tasks.

3. Click Compare Versions to view the differences between the modified values and the original parameter values.
4. Click Run. A message appears, indicating that the action has been submitted.

After the merge task is created, it appears in the list of merge tasks for clusters.

View merge task statistics

On the Business page, choose Business Optimization > File Merging in the left-side navigation pane, and then click the Historical Statistics tab to view the historical statistics of merge tasks for clusters and projects.

Merge task chart

The trend chart for merge tasks displays statistics on the execution of all merge tasks for each day in the last month, including the number of running tasks, number of finished tasks, number of waiting tasks, number of timeout tasks, number of failed tasks, number of invalid tasks, number of merged partitions, number of reduced files, and amount of saved physical storage, in bytes.

Cluster statistics and project statistics

The two tables list statistics on the execution of merge tasks for clusters and projects for a specified day in the last month, including the number of running tasks, number of finished tasks, number of waiting tasks, number of timeout tasks, number of failed tasks, number of invalid tasks, number of merged partitions, number of reduced files, and amount of saved physical storage, in bytes.

3.1.5.2.4.2 File archiving

In the Apsara Bigdata Manager (ABM) console, you can create archive tasks to compress idle files in MaxCompute clusters and projects to save storage space for the clusters.

Definition

ABM sorts the tables or partitions created more than 90 days ago in a cluster by storage space, and then compresses the first 100,000 tables or partitions.

Create an archive task for a cluster

When excessive idle files exist in most projects of a MaxCompute cluster, you can create an archive task to compress the idle files in the cluster in a unified manner.

1. On the Business page, choose Business Optimization > File Archiving in the left-side navigation pane. The Archive Tasks tab appears.
2. In the Archive Tasks for Clusters section, click Create Archive Task, and set relevant parameters in the dialog box that appears.

The following table describes the required parameters.

Parameter	Description
Cluster	The cluster in which you want to run the archive task. Select a cluster from the drop-down list.
Start Time	The start time of the archive task.

Parameter	Description
End Time	The end time of the archive task.
Bandwidth Limit	<p>Specifies whether to support the concurrency of the archive tasks for the cluster.</p> <ul style="list-style-type: none"> • Yes: indicates that archive tasks cannot be run simultaneously. • No: indicates that archive tasks can be run simultaneously.
Maximum Concurrent Jobs	The maximum number of archive tasks that can be run at the same time in the selected cluster. This parameter is valid only when Bandwidth Limit is set to No.
Enabled	Specifies whether the archive task is enabled.
Maximum Running Jobs	The maximum number of jobs that can be run at the same time in the selected cluster. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the selected cluster, not only the archive tasks.
Archive Parameters	<p>The parameter configuration for the archive task. You can use the following default configuration:</p> <pre> { "odps.idata.useragent": "SRE Archive", "odps.oversold.resources.ratio": "100", "odps.merge.quickmerge.flag": "true", "odps.merge.cross.paths": "true", "odps.merge.smallfile.filesize.threshold": "4096", "odps.merge.maxmerged.filesize.threshold": "4096", "odps.merge.max.filenummer.per.instance": "10000", "odps.merge.max.filenummer.per.job": "10000000", "odps.merge.maintain.order.flag": "true", "odps.sql.hive.compatible": "true", "odps.merge.compression.strategy": "normal", "odps.compression.strategy.normal.compressor": "zstd", "odps.merge.failure.handling": "any", "odps.merge.archive.flag": "true" } </pre>

3. Click Compare Versions to view the differences between the modified values and the original parameter values.
4. Click Run. A message appears, indicating that the action has been submitted.

After the archive task is created, it appears in the list of archive tasks for clusters.

Create an archive task for a project

When excessive idle files exist in only a few projects of a MaxCompute cluster, you can create an archive task to compress the idle files in each project.



Note:

If the tables or partitions of a project are not ranked top 100,000 in the cluster of the project, the archive task cannot compress the idle files in the project.

1. On the Business page, choose Business Optimization > File Archiving in the left-side navigation pane. The Archive Tasks tab appears.
2. In the Archive Tasks for Projects section, click Create Archive Task, and set relevant parameters in the dialog box that appears.

The following table describes the required parameters.

Parameter	Description
Region	The region where the cluster of the selected project resides. Select a region from the drop-down list.
Project Name	The name of the project in which you want to run the archive task. Select a project from the drop-down list.
Start Time	The start time of the archive task.
Priority	The priority of the archive task. A smaller value indicates a higher priority.
End Time	The end time of the archive task.
Bandwidth Limit	Specifies whether to support the concurrency of the archive tasks for the project. <ul style="list-style-type: none"> • Yes: indicates that archive tasks cannot be run simultaneously. • No: indicates that archive tasks can be run simultaneously.
Maximum Concurrent Jobs	The maximum number of archive tasks that can be run at the same time in the cluster of the selected project. This parameter is valid only when Bandwidth Limit is set to No.
Enabled	Specifies whether the archive task is enabled.
Maximum Running Jobs	The maximum number of jobs that can be run at the same time in the cluster of the selected project. This parameter is a global parameter. Note that the jobs here refer to all types of jobs in the cluster of the selected project, not only the archive tasks.

3. Click **Compare Versions** to view the differences between the modified values and the original parameter values.
4. Click **Run**. A message appears, indicating that the action has been submitted.

After the archive task is created, it appears in the list of archive tasks for projects

.

View archive task statistics

On the **Business** page, choose **Business Optimization > File Archiving** in the left-side navigation pane, and then click the **Historical Statistics** tab to view the historical statistics of archive tasks for clusters and projects.

Archive task chart

The trend chart for archive tasks displays statistics on the execution of all archive tasks for each day in the last month, including the number of running tasks, number of finished tasks, number of waiting tasks, number of timeout tasks, number of failed tasks, number of invalid tasks, number of merged partitions, number of reduced files, and amount of saved physical storage, in bytes.

Cluster statistics and project statistics

The two tables list statistics on the execution of archive tasks for clusters and projects for a specified day in the last month, including the number of running tasks, number of finished tasks, number of waiting tasks, number of timeout tasks, number of failed tasks, number of invalid tasks, number of merged partitions, number of reduced files, and amount of saved physical storage, in bytes.

3.1.5.2.4.3 Resource analysis

Apsara Bigdata Manager (ABM) allows you to analyze the resources for MaxCompute clusters from multiple dimensions so that you can better understand the data storage in MaxCompute. The dimensions include tables, tasks, execution time, start time, and engines.

Tables

From this dimension, you can view the detailed information about all tables in each project, including the number of partitions, storage space, number of Apsara Distributed File System files, ranking of the number of partitions, and ranking of the number of Apsara Distributed File System files. You can also sort tables based

on the number of partitions, storage space, or number of Apsara Distributed File System files.

On the Business page, choose **Business Optimization > Resource Analysis** in the left-side navigation pane. The **Tables** tab appears.

Projects

From this dimension, you can view the detailed information about storage for each project, including the number of Apsara Distributed File System files, storage space, CU usage, memory usage, number of tasks, number of tables, idle storage, and daily and weekly increases of these items.

On the Business page, choose **Business Optimization > Resource Analysis** in the left-side navigation pane, and then click the **Projects** tab. The **Projects** tab appears.

Tasks

From this dimension, you can view the detailed information about tasks in each project, including the ID of the task instance, running status, CU usage, start time, end time, execution time, ranking of CU usage, and SQL statements.

On the Business page, choose **Business Optimization > Resource Analysis** in the left-side navigation pane, and then click the **Tasks** tab. The **Tasks** tab appears.

Execution time

From this dimension, you can view the numbers of tasks whose execution time is within 5 minutes, within 15 minutes, within 30 minutes, within 60 minutes, and over 60 minutes respectively in each project. The execution time chart displays the trend lines of the numbers of tasks with different execution time by day in different colors.

On the Business page, choose **Business Optimization > Resource Analysis** in the left-side navigation pane, and then click the **Execution Time** tab. The **Execution Time** tab appears.

Start time

From this dimension, you can view the numbers of tasks started in different time periods for each project. The time interval is set to 30 minutes. The tasks chart displays the trend line of the number of tasks started in the specified time period by day.

On the Business page, choose Business Optimization > Resource Analysis in the left-side navigation pane, and then click the Start Time tab. The Start Time tab appears.

Engines

From this dimension, you can view the trend lines of performance statistics of tasks for each project, including CPU usage (`cost_cpu`), memory usage (`cost_mem`), execution time (`cost_time`), input in the unit of bytes (`input_bytes`), input per CU in the unit of bytes (`input_bytes_per_cu`), number of input records (`input_records`), number of input records per CU (`input_records_per_cu`), output in the unit of bytes (`output_bytes`), output per CU in the unit of bytes (`output_bytes_per_cu`), number of output records (`output_records`), and number of output records per CU (`output_records_per_cu`).

On the Business page, choose Business Optimization > Resource Analysis in the left-side navigation pane, and then click the Engines tab. The Engines tab appears.

3.1.5.3 Service O&M

3.1.5.3.1 Control service O&M

3.1.5.3.1.1 Control service O&M overview

This topic describes the features of control service O&M and how to access the control service O&M page.

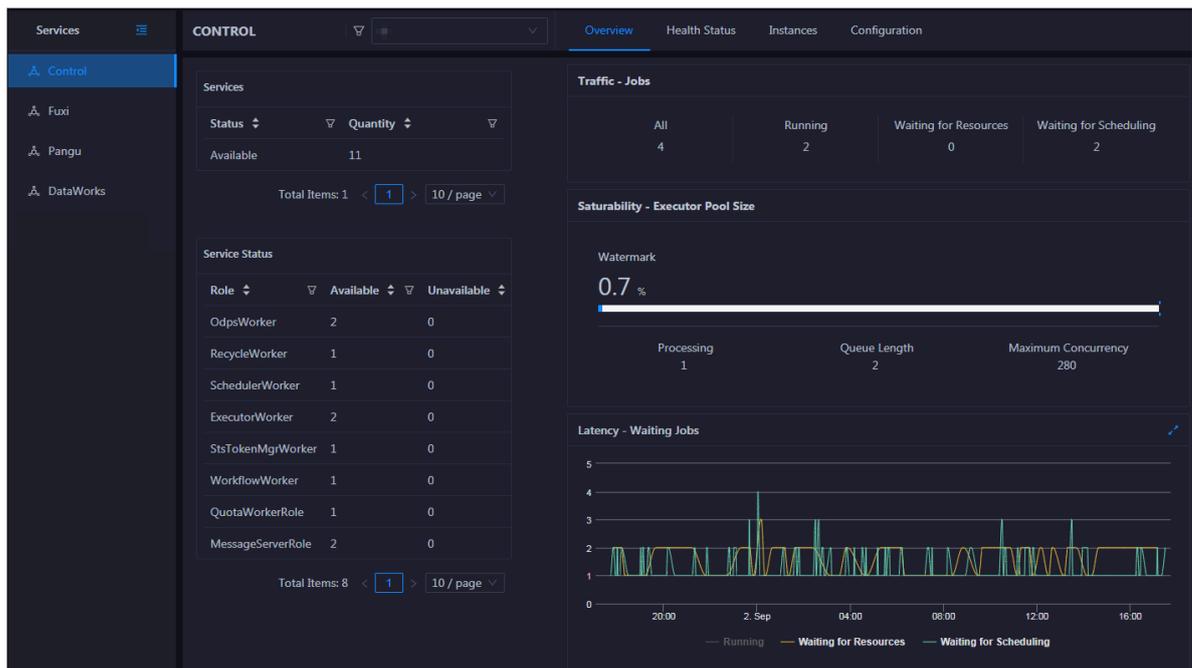
Modules

- **Overview page:** displays the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.
- **Health Status page:** displays all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host .
- **Instances page:** displays information about the service roles, including the host, service status, CPU application, and memory application of each service role.
- **Configuration page:** provides the access to configuring global computing, cluster-level computing, computing scheduling, and cluster endpoints.

- **Start Service Role or Stop Service Role action:** allows you to enable or disable the control service roles of MaxCompute and view the execution history. You can also locate the failure cause when service role disabling or enabling fails.
- **Start Admin Console action:** allows you to start AdminConsole.
- **Collect Service Logs action:** allows you to collect the service logs for the specified time period. This helps you locate the failure cause.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Control in the left-side navigation pane. The Overview page for the control service appears.

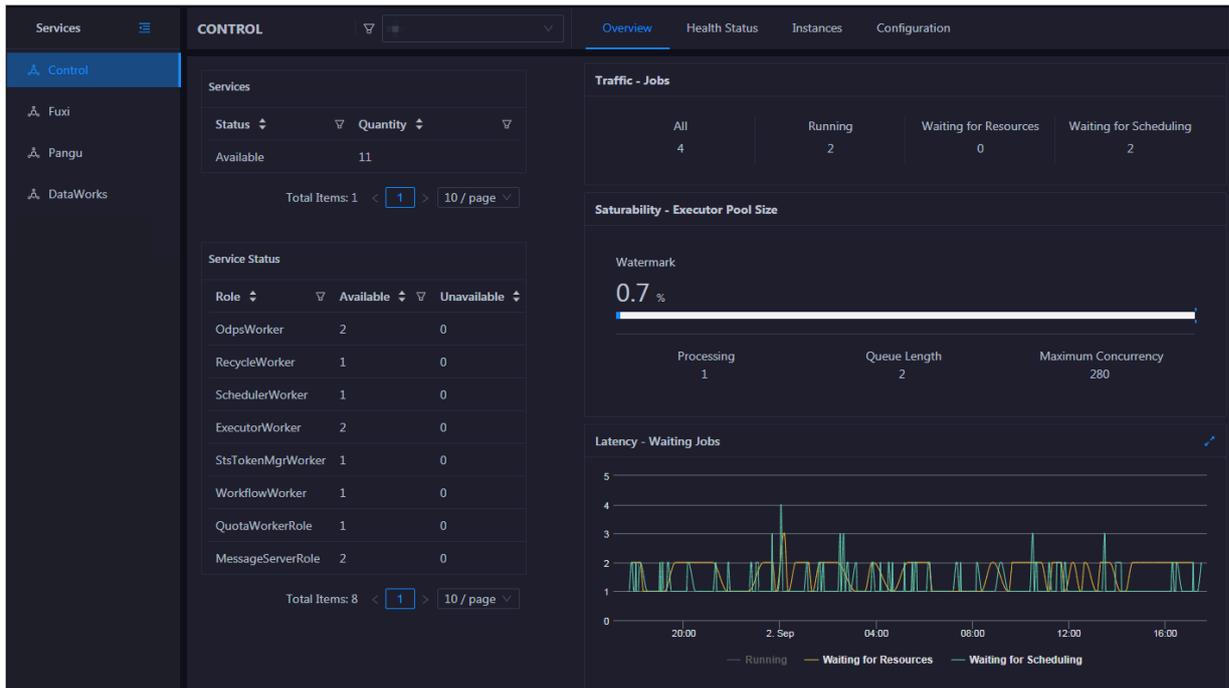


3.1.5.3.1.2 Control service overview

The Overview page displays the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.

Entry

On the Services page, click Control in the left-side navigation pane. The Overview page for the control service appears.



On the Overview page, you can view the overall running information about the control service, including the service summary, service status, job summary, executor pool summary, and job status.

Services

This section displays the numbers of available services and unavailable services respectively.

Service Status

This section displays all control service roles. You can also view the numbers of available and unavailable services respectively for each service role.

Traffic - Jobs

This section displays the total number of jobs in the cluster, and the numbers of running jobs, jobs waiting for resources, and jobs waiting for scheduling respectively.

Saturability - Executor Pool Size

The section displays information about the thread pool, including the resource usage, number of jobs being processed, queue length, and maximum concurrency.

Latency - Waiting Jobs

This section displays the trend chart of jobs. The chart displays of the trend lines of the numbers of running jobs, jobs waiting for resources, and jobs waiting for scheduling in different colors.

3.1.5.3.1.3 Control service health

On the Health Status page for the control service, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

On the Services page, click **Control** in the left-side navigation pane, and then click the **Health Status** tab.

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

Other operations

On the Health Status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host. For more information, see [Cluster health](#).

3.1.5.3.1.4 Control service instances

The Instances page displays information about the service roles, including the host, service status, CPU application, and memory application of each service role.

Entry

On the Services page, click **Control** in the left-side navigation pane, and then click the **Instances** tab.

The Instance page displays information about the service roles, including the host, service status, CPU application, and memory application of each service role.

Other operations

You can filter or sort service roles by column to facilitate information display. For more information, see [Common operations](#).

3.1.5.3.1.5 Disable or enable a control service role

Apsara Bigdata Manager (ABM) allows you to disable or enable control service roles of MaxCompute and view the execution history. You can also locate the failure cause when service role disabling or enabling fails.

Disable a service role

1. [Log on to the ABM console](#).
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Control in the left-side navigation pane. Click Actions in the upper-left corner, and then click Stop Service Role.
5. In the dialog box that appears, select a service role to be disabled, and then click Run. A message appears, indicating that the action has been submitted.
6. Click Actions in the upper-left corner, and then click Execution History next to Stop Service Role to check whether the action is successful in the execution history.

The current status, submission time, start time, end time, and operator of each action are recorded in the execution history.

7. Click Details to view the execution details.

On the execution details page, you can view the job name, execution status, execution steps, script, and parameter settings. You can also download the execution details to your local device.

Enable a service role

1. [Log on to the ABM console](#).
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.

4. On the Services page, click Control in the left-side navigation pane. Click Actions in the upper-left corner, and then click Start Service Role.
5. In the dialog box that appears, select a service role to be enabled, and then click Run. A message appears, indicating that the action has been submitted.
6. Click Actions in the upper-left corner, and then click Execution History next to Start Service Role to check whether the action is successful in the execution history.

The current status, submission time, start time, end time, and operator of each action are recorded in the execution history.

7. Click Details to view the execution details.

On the execution details page, you can view the job name, execution status, execution steps, script, and parameter settings. You can also download the execution details to your local device.

Locate the failure cause

The following content uses service role enabling as an example to describe how to locate the failure cause:

1. In the execution history dialog box, click Details in the Details column of the task to view the details.
2. On the page that appears, click View Details for a failed step to locate the failure cause.

You can also view the parameter settings, outputs, error messages, script, and execution parameters to locate the failure cause.

3.1.5.3.1.6 Start AdminConsole

AdminConsole is a management platform of MaxCompute. It is disabled by default. Apsara Bigdata Manager (ABM) allows you to quickly start AdminConsole to better manage MaxCompute clusters.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on MaxCompute.

Step 1: Start AdminConsole

1. [Log on to the ABM console.](#)

2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Control in the left-side navigation pane.
5. On the Control page, click Actions in the upper-left corner, and then click Start Admin Console.
6. In Start Admin Console dialog box that appears, click Run. A message appears, indicating that the action has been submitted.

Step 2: View the execution status or progress

1. On the Control page, click Actions in the upper-left corner, and then click Execution History next to Start Admin Console to view the execution history.
In the Current Status column, RUNNING indicates that the execution is in progress, FAILED indicates that the execution fails, and SUCCESS indicates that the execution is successful.
2. You can click Details in the Details column of a task in the RUNNING state to view the execution progress.

(Optional) Step 3: Locate the failure cause

If the status of the task is FAILED, you can view the execution logs to locate the failure cause.

1. On the Control page, click Actions in the upper-left corner, and then click Execution History next to Start Admin Console to view the execution history.
2. In the execution history dialog box, click Details in the Details column of the task to view the details.
3. On the Servers tab of the failed step, click View Details in the Actions column of a failed server. The Execution Output tab appears in the Execution Details section. You can view the output to locate the failure cause.

3.1.5.3.1.7 Collect service logs

Apsara Bigdata Manager (ABM) allows you to collect the service logs for the specified time period. This helps you locate the failure cause.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on MaxCompute.

Step 1: Collect service logs

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Control in the left-side navigation pane.
5. On the Control page, click Actions in the upper-left corner, and then click Collect Service Logs.
6. In the Collect Service Logs dialog box that appears, set the parameters.

The following table describes the required parameters.

Parameter	Description
Target Service	The target service from which you want to collect service logs. Select a target service from the drop-down list.
Time Period	The time range in which the job logs that you want to collect are generated.
Degree of Concurrency	The maximum number of nodes from which you can collect service logs at the same time.
Hostname	The name of the host. Separate multiple hostnames with commas (,).

7. Click Run. A message appears, indicating that the action has been submitted.

Step 2: View the execution status or progress

1. On the Control page, click Actions in the upper-left corner, and then click Execution History next to Collect Service Logs to view the execution history.

In the Current Status column, **RUNNING** indicates that the execution is in progress, **FAILED** indicates that the execution fails, and **SUCCESS** indicates that the execution is successful.
2. You can click Details in the Details column of a task in the **RUNNING** state to view the execution progress.

(Optional) Step 3: Locate the failure cause

If the status of the task is FAILED, you can view the execution logs to locate the failure cause.

- 1. On the Control page, click Actions in the upper-left corner, and then click Execution History next to Collect Service Logs to view the execution history.**
- 2. In the execution history dialog box, click Details in the Details column of the task to view the details.**
- 3. On the Servers tab of the failed step, click View Details in the Actions column of a failed server. The Execution Output tab appears in the Execution Details section. You can view the output to locate the failure cause.**

3.1.5.3.2 Job Scheduler O&M

3.1.5.3.2.1 Job Scheduler O&M overview

This topic describes the O&M features of Job Scheduler and how to access the O&M page.

Modules

- **Overview page: displays the key operation metrics of Job Scheduler, including the service overview, service status, resource usage, and compute node overview . You can also view the trend charts of CPU and memory usage on this page.**
- **Health Status page: displays the check results of all checkers for Job Scheduler. The check results are divided into Critical, Warning, Exception, and OK.**
- **Quotas page: allows you to view, add, or modify Job Scheduler quota groups.**
- **Instances page: displays information about the service roles of Job Scheduler.**
- **Compute Nodes page: displays all Job Scheduler compute nodes and allows you to add compute nodes to or remove them from the blacklist or read-only list.**
- **Enable SQL Acceleration or Disable SQL Acceleration action: allows you to enable or disable SQL acceleration for Job Scheduler.**
- **Restart Fuxi Master Node action: allows you to restart the primary master node of Job Scheduler.**

Entry

- 1. [Log on to the ABM console.](#)**
- 2. Click  in the upper-left corner, and then click MaxCompute.**

3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Fuxi in the left-side navigation pane and then select a cluster. The Overview page for the selected cluster appears.

3.1.5.3.2.2 Job Scheduler overview

The Overview page displays the key operation metrics of Job Scheduler, including the service overview, service status, resource usage, and compute node overview. You can also view the trend charts of CPU and memory usage on this page.

Entry

1. On the Services page, click Fuxi in the left-side navigation pane.
2. Select a cluster, and then click the Overview tab. The Overview page for the selected cluster appears.

The Overview page displays the key operation metrics of Job Scheduler, including the service overview, service status, resource usage, and compute node overview. You can also view the trend charts of CPU and memory usage on this page.

Services

This section displays the numbers of available services, unavailable services, and services that are being upgraded respectively.

Roles

This section displays all Job Scheduler service roles and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

Saturability - Resource Usage

This section displays the usage and allocation of CPU and memory resources.

- **CPU (Core):** displays the CPU usage, the total number of CPU cores, the number of available CPU cores, and the CPU cores for SQL acceleration.
- **Memory (Bytes):** displays the memory usage, the total memory size, the available memory size, and the size of memory for SQL acceleration.

CPU Usage (1/100 Core) and Memory Usage (MB)

This section displays trend charts of CPU and memory usage for Job Scheduler . Each trend chart displays the trend lines of the used quota, minimum quota , maximum cluster quota, requested quota, and maximum quota over time in different colors.

Click  in the upper-right corner of a chart to zoom in it. The following figure shows an enlarged CPU usage chart.

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

Compute Nodes

This section displays the details of Job Scheduler compute nodes, including the online rate, the total number of compute nodes, the number of online compute nodes, and the number of compute nodes in the blacklist.

3.1.5.3.2.3 Job Scheduler health

On the Health Status page for Job Scheduler, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. On the Services page, click Fuxi in the left-side navigation pane.
2. Select a cluster, and then click the Health Status tab. The Health Status page for Job Scheduler appears.

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

Other operations

On the Health Status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear

alerts (if any). In addition, you can log on to a host and perform manual checks on the host. For more information, see [Cluster health](#).

3.1.5.3.2.4 Job Scheduler quota management

You can view, add, or modify Job Scheduler quota groups on the Quotas page. A quota group is used to allocate computing resources, including CPU and memory resources, to MaxCompute projects.

Entry

1. On the Services page, click Fuxi in the left-side navigation pane.
2. Select a cluster, and then click the Quotas tab. The Quotas page for the selected cluster appears.

The Quotas page lists the existing Job Scheduler quota groups.

Add a quota group

1. On the Quotas page, click Create Quota Group in the upper-left corner.
2. In the Quota Group dialog box that appears, set the parameters as instructed.
3. Click Run. A message appears, indicating that the action has been submitted.

After the quota group is created, it appears in the quota group list.

View quota group details

Click the name of a quota group to view the details. The Resource Usage tab displays the trend charts of CPU and memory usage. The Applications page displays the projects that use the quota group resources.

Figure 3-2: Resource usage

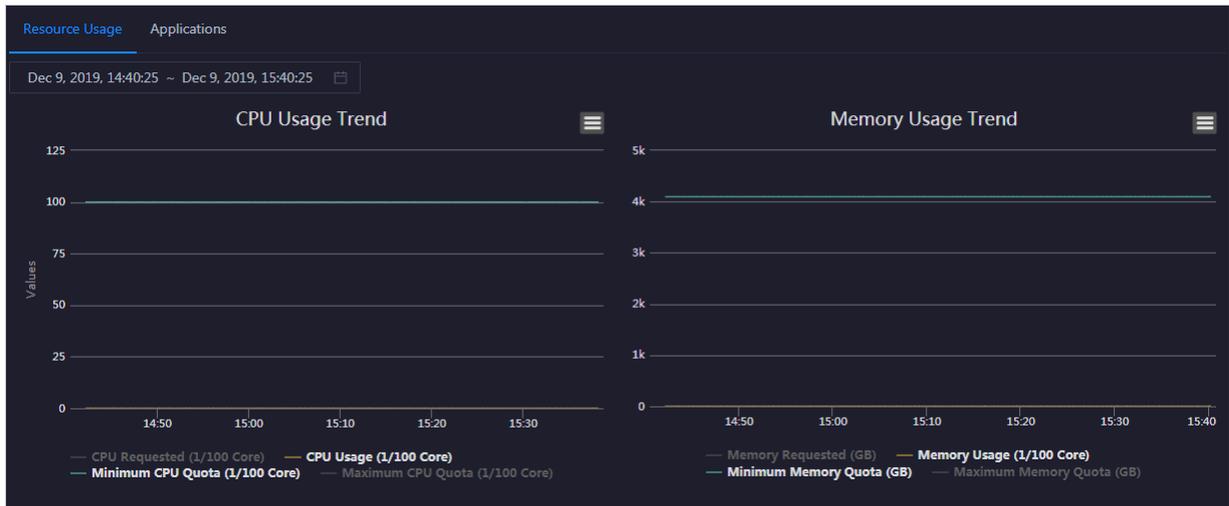


Figure 3-3: Applications

The screenshot shows the 'Applications' tab with a table of projects. The table has columns for Project, owner, BU, Created At, and Description. The data row shows a project with owner 'ALYUN\$', BU 'Default', and Created At '2019-10-28 03:21:50'. Below the table, there is a pagination control showing 'Total Items: 1', a page number '1' in a box, '10 / page', and a 'Goto' button.

Project	owner	BU	Created At	Description
	ALYUN\$	Default	2019-10-28 03:21:50	

Modify a quota group

- 1. On the Quotas page, find the quota group you want to modify, click Modify in the Actions column, and then modify parameters as instructed in the dialog box that appears.**
- 2. Click Run. A message appears, indicating that the action has been submitted. You can check whether the quota group is successfully modified in the quota group list after the configuration is completed.**

3.1.5.3.2.5 Job Scheduler instances

The Instances page displays information about the Job Scheduler service roles, including the service role name, service role host, service role status, and host status.

Entry

1. On the Services page, click **Fuxi** in the left-side navigation pane.
2. Select a cluster, and then click the **Instances** tab. The Instances page for Job Scheduler appears.

On the Instances page, you can view information about the Job Scheduler service roles, including the service role name, service role host, service role status, and host status.

Other operations

You can filter or sort service roles by column to facilitate information display. For more information, see [Common operations](#).

3.1.5.3.2.6 Job Scheduler compute nodes

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active. You can also add compute nodes to or remove them from the blacklist or read-only list on the Compute Nodes page.

Entry

1. On the Services page, click **Fuxi** in the left-side navigation pane.
2. Select a cluster, and then click the **Compute Nodes** tab. The Compute Nodes page for Job Scheduler appears.

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active.

Blacklist and read-only setting

You can add compute nodes to or remove them from the blacklist or read-only list. To add compute nodes to the blacklist, follow these steps:

1. On the **Compute Nodes** page, click **Actions** for the target compute node and then select **Add to Blacklist**.
2. In the dialog box that appears, enter the hostname. If you want to add multiple compute nodes to the blacklist, separate hostnames with commas (,).
3. Click **Run**. A message appears, indicating that the action has been submitted.

You can view the blacklist statuses of compute nodes in the compute node list after the configuration is completed.

3.1.5.3.2.7 Enable or disable SQL acceleration

You can enable or disable SQL acceleration for Job Scheduler in the Apsara Bigdata Manager (ABM) console. Enabling SQL acceleration can greatly increase the speed of running SQL statements in Job Scheduler, but it consumes more computing resources.

Enable SQL acceleration

1. On the **Services** page, click **Fuxi** in the left-side navigation pane and then select a cluster.
2. Click **Actions** next to **FUXI** in the upper-left corner, and then click **Enable SQL Acceleration**.
3. In the dialog box that appears, set **WorkerSpans**.

WorkerSpans: the default resource quota of the cluster and the resource quota within the specified period. Default value: default:2,12-23:2. The default value indicates that the default resource quota is 2 and the resource quota for the period from 12:00 to 23:00 is also 2. You can set the resource quota as needed. For example, during business peak hours, you can set this parameter to default:2,12-23:4 to increase the resource quota.

4. Click **Run**. A message appears, indicating that the action has been submitted.

Disable SQL acceleration

1. On the **Services** page, click **Fuxi** in the left-side navigation pane and then select a cluster.
2. Click **Actions** next to **FUXI** in the upper-left corner, and then click **Disable SQL Acceleration**.
3. In the dialog box that appears, click **Run**. A message appears, indicating that the action has been submitted.

View the execution history of enabling or disabling SQL acceleration

After you submit the action of enabling or disabling SQL acceleration, you can check whether the current action is completed by viewing the execution history. The system executes the action as a job. It provides execution records and logs for each execution so that you can locate faults encountered during the execution of the job. To view the execution history of enabling SQL acceleration, follow these steps:

1. On the Services page, click Fuxi in the left-side navigation pane and then select a cluster.
2. Click Actions next to FUXI in the upper-left corner, and then click  next to Enable SQL Acceleration.
3. In the dialog box that appears, view the execution history of enabling SQL acceleration.

The current status, submission time, start time, end time, and operator of each action are recorded in the execution history.

4. If the execution fails, click Details to go to the Jobs page to locate the failure cause.

3.1.5.3.2.8 Restart the primary master node of Job Scheduler

Job Scheduler is the resource management and task scheduling system of the Apsara system. Apsara Bigdata Manager (ABM) allows you to quickly restart the primary master node of Job Scheduler. Cluster services are not affected during the restart process.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on MaxCompute.

Step 1: Restart the primary master node of Job Scheduler

1. [Log on to the ABM console](#).
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Fuxi in the left-side navigation pane.

5. On the page that appears, click **Actions** in the upper-left corner, and then click **Restart Fuxi Master Node**.
6. In the **Restart Fuxi Master Node** dialog box that appears, click **Run**. A message appears, indicating that the action has been submitted.

Step 2: View the execution status or progress

1. On the **Fuxi** page, click **Actions** in the upper-left corner, and then click **Execution History** next to **Restart Fuxi Master Node** to view the execution history.

In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **FAILED** indicates that the execution fails, and **SUCCESS** indicates that the execution is successful.

2. You can click **Details** in the **Details** column of a task in the **RUNNING** state to view the execution progress.

(Optional) Step 3: Locate the failure cause

If the status of the task is **FAILED**, you can view the execution logs to locate the failure cause.

1. On the **Fuxi** page, click **Actions** in the upper-left corner, and then click **Execution History** next to **Restart Fuxi Master Node** to view the execution history.
2. In the execution history dialog box, click **Details** in the **Details** column of the task to view the details.
3. On the **Servers** tab of the failed step, click **View Details** in the **Actions** column of a failed server. The **Execution Output** tab appears in the **Execution Details** section. You can view the output to locate the failure cause.

3.1.5.3.3 Apsara Distribute File System O&M

3.1.5.3.3.1 Apsara Distribute File System O&M overview

This topic describes the O&M features of Apsara Distributed File System and how to access the O&M page.

Modules

- **Overview page:** displays the key operation metrics of Apsara Distributed File System, including the service overview, service status, storage usage, and storage node overview. You can also view the trend charts of storage usage and file count on this page.

- **Health Status page:** displays the check results of all checkers for Apsara Distributed File System. The check results include Critical, Warning, Exception, and OK.
- **Instances page:** displays information about the service roles of Apsara Distributed File System.
- **Storage Nodes page:** displays information about the storage nodes of Apsara Distributed File System. On this page, you can set the status of a storage node to Disabled or Normal. In addition, you can set the status of a disk on a storage node to Normal or Error.
- **Change Primary Master Node action:** allows you to change the primary master node of Apsara Distributed File System in a cluster.
- **Empty Recycle Bin action:** allows you to empty the recycle bin of Apsara Distributed File System.
- **Enable Data Rebalancing or Disable Data Rebalancing action:** allows you to enable or disable the data rebalancing feature of Apsara Distributed File System.
- **Run Checkpoint on Master Node action:** allows you to run checkpoints on master nodes of Apsara Distributed File System to write memory data into disks.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Pangu in the left-side navigation pane and then select a cluster. The Overview page for the selected cluster appears.

3.1.5.3.3.2 Apsara Distributed File System overview

On the Overview page for Apsara Distributed File System, you can view the key operation metrics, including the service overview, service status, storage usage, and storage node overview. You can also view the trend charts of storage usage and file count on this page.

Entry

1. On the Services page, click Pangu in the left-side navigation pane.

2. Select a cluster, and then click the Overview tab. The Overview page for the selected cluster appears.

The Overview page displays the key operation metrics of Apsara Distributed File System, including the service overview, service status, health check result, health check history, storage usage, and storage nodes. You can also view the trend charts of storage usage and file count on this page.

Services

This section displays the status of Apsara Distributed File System and the number of service roles.

Roles

This section displays all Apsara Distributed File System service roles and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

Saturability - Storage

This section displays the storage usage (Storage) and file count (File Count).

- **Storage:** displays the storage usage, total storage size, available storage size, and recycle bin size.
- **File Count:** displays the file count usage, maximum number of files, number of existing files, and number of files in the recycle bin.

Storage Trend and File Count Trend

This section displays the storage usage and file count charts. The storage usage chart displays the trend lines of the total storage size, used storage size, and storage usage over time in different colors. The file count chart displays the trend line of the file count.

Click  in the upper-right corner of a chart to zoom in it. The following figure shows an enlarged storage usage chart.

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

Storage Nodes

This section displays information about the storage nodes of Apsara Distributed File System, including the respective number of data nodes, normal nodes, disks

, and normal disks. You can also view the faulty node percentage and faulty disk percentage in this section.

3.1.5.3.3.3 Service instances

The Instances page displays information about the Apsara Distributed File System service roles, including the service role name, service role host, service role status, and host status.

Entry

1. On the Services page, click Pangu in the left-side navigation pane.
2. Select a cluster, and then click the Instances tab. The Instances page for Apsara Distributed File System appears.

On the Instances page, you can view information about the Apsara Distributed File System service roles, including the service role name, service role host, service role status, and host status.

Other operations

You can filter or sort service roles by column to facilitate information display. For more information, see [Common operations](#).

3.1.5.3.3.4 Service health

On the Health Status page for Apsara Distributed File System, you can view all health check items, check details, check results, and solutions for alerts (if there are any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. On the Services page, click Pangu in the left-side navigation pane.
2. Select a cluster, and then click the Health Status tab. The Health Status page for Apsara Distributed File System appears.

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

Other operations

On the Health Status page, you can view checker details, hosts with alerts, and alert causes. You can also log on to hosts with alerts, clear alerts, and run checkers again. For more information, see [Cluster health](#).

3.1.5.3.3.5 Apsara Distributed File System storage nodes

The Storage Nodes page displays the information about the storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, time to live (TTL), and send buffer size. You can also set the status of storage nodes and data disks on this page.

Entry

- 1. On the Services page, click Pangu in the left-side navigation pane.**
- 2. Select a cluster, and then click the Storage Nodes tab. The Storage Nodes page for Apsara Distributed File System appears.**

On the Storage Nodes page, you can view storage node details, including the total CPU, available CPU, total storage size, and available storage size. You can also check whether a node is added to the blacklist and whether it is active.

Set the storage node status

You can set the storage node status to Disabled or Normal. This section describes how to set the status of a storage node to Disabled.

- 1. On the Storage Nodes page, click Actions for the target storage node, and then select Set Node Status to Disabled.**
- 2. In the dialog box that appears, set volume and hostname. If you need to add multiple storage nodes to the blacklist, separate the storage node names with commas (,).**
- 3. Click Run. A message appears, indicating that the action has been submitted.**

You can view the status of storage nodes in the storage node list.

Set the data disk status

You can set the data disk status to Error or Normal. This section describes how to set the status of a data disk to Error.

- 1. On the Storage Nodes page, click Actions for the target data disk, and then select Set Disk Status to Error.**

2. In the dialog box that appears, set volume, hostname, and diskid. If you need to add multiple data disks to the blacklist, separate the disk names with commas (,).
3. Click Run. A message appears, indicating that the action has been submitted.

3.1.5.3.3.6 Change the primary master node for Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to perform primary/secondary switchover on the master nodes of Apsara Distributed File System. After the primary/secondary switchover is completed, a secondary master node becomes the new primary master node, and the original primary master node becomes a new secondary master node.

Prerequisites

- Your ABM account must have the required permissions to perform O&M operations on MaxCompute.
- You have obtained the roles of the primary and secondary master nodes in a volume. To view the role of a master node, log on to the Apsara Infrastructure Management Framework console and access the PanguTools# host in the MaxCompute cluster. Then, run the `puadmin gems` command on the host.
- You have obtained the hostname of the secondary master node that is to be the new primary master node. Log on to the ABM console, go to the MaxCompute O&M page, and then click Services. On the page that appears, click Pangu in the left-side navigation pane, and then click the Instances tab. On the Instances page, view the hostnames of PanguMaster# hosts.

Context

A volume in Apsara Distributed File System is similar to a namespace in Hadoop Distributed File System (HDFS). The default volume is PanguDefaultVolume. Multiple volumes may exist if a cluster consists of numerous nodes. A volume has three master nodes. One of the nodes serves as the primary master node, whereas the other two nodes serve as secondary master nodes.

Procedure

1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.
2. Click Actions, and then select Change Primary Master Node.

3. In the dialog box that appears, set relevant parameters.

The parameters are described as follows:

- **Volume:** the volume whose primary master node needs to be changed. Default value: PanguDefaultVolume. If a cluster consists of multiple volumes, set this parameter to the name of the actual volume whose primary master node needs to be changed.
- **Hostname:** the hostname of the secondary master node that is to be the new primary master node.
- **log_gap:** the maximum log number gap between the original primary and secondary master nodes. During the switchover, the system checks the log number gap between the original primary and secondary master nodes. If the gap is less than the specified value, switchover is allowed. Otherwise, you cannot change the primary master node. Default value: 100000.

4. Click Run. A message appears, indicating that the action has been submitted.

5. View the execution status.

Click Actions, and then click  next to Change Primary Master Node to view the execution history.

In the Current Status column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

6. Click Details for a failed execution to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

3.1.5.3.3.7 Empty the recycle bin of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to empty the recycle bin of Apsara Distributed File System.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on MaxCompute.

Procedure

1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.
2. Click Actions, and then select Empty Recycle Bin.
3. In the dialog box that appears, set volume. The default value is PanguDefaultVolume.
4. Click OK. A message is displayed, indicating that the action has been submitted.
5. View the execution status.

Click Actions, and then click  next to Empty Recycle Bin to view the execution history.

In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

6. Click Details for a failed execution to locate the failure cause.

You can also view information about parameter settings, hosts, script, and execution parameters to locate the failure cause.

3.1.5.3.3.8 Enable or disable data rebalancing for Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to enable or disable data rebalancing for Apsara Distributed File System.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on MaxCompute.

Disable data rebalancing

1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.
2. Click Actions, and then select Disable Data Rebalancing.
3. In the dialog box that appears, set volume. The default value is PanguDefaultVolume.
4. Click OK. A message is displayed, indicating that the action has been submitted.

5. View the execution status.

Click **Actions**, and then click **Execution History** next to **Disable Data Rebalancing** to view the execution history.

In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

6. Click Details for a failed execution to locate the failure cause. For more information, see [Locate the failure cause](#).

Enable data rebalancing

1. On the **Services** page, click **Pangu** in the left-side navigation pane, and then select a cluster. The **Overview** page for **Apsara Distributed File System** appears.
2. Click **Actions**, and then select **Enable Data Rebalancing**.
3. In the dialog box that appears, set volume. The default value is **PanguDefaultVolume**.
4. Click **OK**. A message is displayed, indicating that the action has been submitted.
5. View the execution status.

Click **Actions**, and then click  next to **Enable Data Rebalancing** to view the execution history.

In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

6. Click Details for a failed execution to locate the failure cause. For more information, see [Locate the failure cause](#).

Locate the failure cause

The following describe how to locate the failure cause for enabling data rebalancing.

1. In the **Enable Data Rebalancing Execution History** dialog box, click **Details** a failed execution.

2. On the page that appears, click **View Details** for a failed step to locate the failure cause.

You can also view information about parameter settings, hosts, script, and execution parameters to locate the failure cause.

3.1.5.3.3.9 Run a checkpoint on the master nodes of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to run checkpoints on master nodes of Apsara Distributed File System. This operation writes memory data into disks.

When a failure occurs to Apsara Distributed File System, you can use checkpoints to restore data to the status before the failure. This ensures data consistency.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on MaxCompute.

Procedure

1. On the **Services** page, click **Pangu** in the left-side navigation pane, and then select a cluster. The **Overview** page for Apsara Distributed File System appears.
2. Click **Actions**, and then select **Run Checkpoint on Master Node**.
3. In the dialog box that appears, set volume. The default value is **PanguDefaultVolume**.
4. Click **OK**. A message is displayed, indicating that the action has been submitted.
5. View the execution status.

Click **Actions**, and then click  next to **Run Checkpoint on Master Node** to view the execution history.

In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

6. Click **Details** for a failed execution to locate the failure cause.

You can also view information about the parameter settings, hosts, script, and execution parameters to locate the failure cause.

3.1.5.3.4 DataWorks O&M

3.1.5.3.4.1 DataWorks O&M overview

This topic describes the O&M features of DataWorks and how to access the DataWorks O&M page.

Modules

- **Overview page:** displays the key operation metrics, including service overview, service status, instance scheduling information, and slot usage. On the page, you can also view the trend chart of finished tasks.
- **Health Status page:** displays the check results of all checkers for DataWorks. The check results are divided into Critical, Warning, Exception, and OK.
- **Instances page:** displays the service roles of DataWorks.
- **Slots page:** displays the information about slot usage in DataWorks and allows you to modify the number of slots in resource groups and hosts.
- **Tasks page:** displays the running status of DataWorks tasks.
- **Add Server to Scale out Cluster or Remover Server to Scale in Cluster action:** allows you to scale in or out a DataWorks cluster.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click DataWorks in the left-side navigation pane. The Overview page for DataWorks appears.

3.1.5.3.4.2 DataWorks overview

The DataWorks overview page displays the key operation metrics, including service overview, service status, instance scheduling information, slot usage, and finished tasks.

Entry

On the Services page, click DataWorks in the left-side navigation pane. The Overview page for DataWorks appears.

The Overview page displays the key operation metrics, including service overview, service status, instance scheduling information, and slot usage. On the page, you can also view the trend chart of finished tasks.

Services

This section displays the numbers of available services, unavailable services, and services that are being upgraded respectively.

Roles

This section displays all DataWorks service roles and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

Scheduler

This section displays the number of successful instances, number of instances not running, waiting time, number of running instances, number of failed instances, and number of instances waiting for resources.

Saturability - Slot Usage

This section displays the total number of slots, the number of used slots, the number of unavailable slots, and the number of idle slots for DataWorks. Slots are resources that can be used by DataWorks for instance scheduling.

Finished Tasks

This section displays the trend chart of finished tasks. The trend chart displays the trend lines of the number of tasks finished yesterday, the number of tasks finished today, and the average number of finished tasks in different colors.

3.1.5.3.4.3 DataWorks health

On the Health Status page for DataWorks, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

On the Services page, click DataWorks in the left-side navigation pane, and then click the Health Status tab. The Health Status page for DataWorks appears.

The Health Status page displays all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the

Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

Other operations

On the Health Status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host. For more information, see [Cluster health](#).

3.1.5.3.4.4 DataWorks instances

The Instances page under DataWorks displays information about all DataWorks service roles, including the name, status, expected, and actual numbers of machines in the final status, and number of machines in cleanup.

Entry

On the Services page, click DataWorks in the left-side navigation pane, and then click the Instances tab. The Instances page for DataWorks appears.

The Instances page displays information about all DataWorks service roles, including the status and the expected and actual numbers of machines in the final status. The statuses include good, bad, and upgrading.

Other operations

You can filter or sort service roles by column to facilitate information display. For more information, see [Common operations](#).

3.1.5.3.4.5 DataWorks slots

Slots are resources used to process tasks. Apsara Bigdata Manager (ABM) allows you to view the slot information of DataWorks clusters, resource groups, and hosts, including the maximum number of slots, the number of used slots, and the slot usage. You can also migrate resource groups, modify the number of slots for resource groups or hosts, and modify the host status.

Concepts

A data migration unit (DMU) represents the minimum operating capability required by a Data Integration task, that is, the data synchronization processing capability given limited CPU, memory, and network resources.

Resources measured by DMU are allocated by slot. Each DMU occupies two slots.

Entry

On the Services page, click DataWorks in the left-side navigation pane, and then click the Slots tab. The Slots page for DataWorks appears.

The Slots page consists of the Cluster, Groups, and Host tabs for you to view the slot information of the clusters, resource groups, and hosts.

Cluster slots

The Cluster page displays the slot overview of all DataWorks clusters, including the total number of slots, the numbers of used slots and available slots, and the slot usage. The page also displays the cluster running status.

For more information about slots of a specified cluster, click the name of the cluster .

On the cluster details page, you can view the numbers of gateways, resource groups , slots, used slots, and available slots, and the slot usage of the cluster at the top. You can also view the trend chart of slot usage over time at the bottom. The trend chart displays the trend lines of the number of used slots, the maximum number of slots, and the number of available slots in different colors.

You can click the name of a metric under the chart to determine whether to display the corresponding trend line in the chart. A highlighted metric name indicates that the corresponding trend line is displayed, whereas a dimmed metric name indicates that the corresponding trend line is not displayed.

Resource group slots

The Groups page displays the slot overview of all DataWorks resource groups, including the maximum number of slots, the numbers of used slots and available slots, and the slot usage. The page also displays the name, cluster, project, and running status of each resource group.

To view more information about slots of a specified resource group, click the ID of the resource group.

On the resource group details page, you can view the current slot information of the resource group, for example, the number of used slots and the maximum number of slots, at the top. You can also view the trend chart of slot usage over time, the nodes that occupy the slots, and the owners at the bottom. The trend chart displays

the trend lines of the number of used slots, the maximum number of slots, and the number of available slots in different colors.

You can click the name of a metric under the chart to determine whether to display the corresponding trend line in the chart. A highlighted metric name indicates that the corresponding trend line is displayed, whereas a dimmed metric name indicates that the corresponding trend line is not displayed.

Modify the number of resource group slots

If the number of slots in a resource group is insufficient or excessive, you can modify the number of slots to add or remove resources in the resource group.

- 1. On the Groups page, find the target resource group, and click Modify Slots in the Actions column.**
- 2. In the dialog box that appears, set Maximum Slots.**
- 3. Click Run. A message appears, indicating that the action has been submitted.**

Migrate a resource group

If the slots in a cluster bound to a resource group are insufficient and cannot be increased, you can bind the resource group to another cluster.

- 1. On the Groups page, find the target resource group, and click Migrate Resource Group in the Actions column.**
- 2. In the dialog box that appears, set Target Cluster.**
- 3. Click Run. A message appears, indicating that the action has been submitted.**

Host slots

The Host page displays the slot overview of all DataWorks hosts, including the maximum number of slots, the number of used slots, and the slot usage. The page also displays the IP address, cluster, running status, activeness, and monitoring status of each host.

To view more information about slots of a specified host, click the name of the host.

On the host details page, you can view the current slot information of the host, for example, the number of used slots and the maximum number of slots, at the top. You can view the trend chart of slot usage over time at the bottom. The trend chart displays the trend lines of the number of used slots, the maximum number of slots, and the number of available slots in different colors.

You can click the name of a metric under the chart to determine whether to display the corresponding trend line in the chart. A highlighted metric name indicates that the corresponding trend line is displayed, whereas a dimmed metric name indicates that the corresponding trend line is not displayed.

Modify the host status

If the number of slots in a host is insufficient or excessive, you can modify the number of slots to add or remove resources in the host.

1. On the Host page, find the target host and click **Modify Status**.
2. In the dialog box that appears, set **Status**.
3. Click **Run**. A message appears, indicating that the action has been submitted.

Modify the number of host slots

The host can be in the normal, unavailable, or suspended state. You can modify the host status as needed.

1. On the Host page, find the target host and click **Modify Slots**.
2. In the dialog box that appears, set **Maximum Slots**.
3. Click **Run**. A message appears, indicating that the action has been submitted.

3.1.5.3.4.6 DataWorks tasks

The **Tasks** page under **DataWorks** displays tasks created by a user in **DataWorks**. You can filter or sort tasks by column to facilitate information display.

Entry

1. On the **Services** page, click **DataWorks** in the left-side navigation pane.
2. Select a cluster, and then click the **Tasks** tab. The **Tasks** page for **DataWorks** appears.

The **Tasks** page displays the task information of the current cluster, including the project name, node name, node ID, business date, owner, running status, start time, end time, priority, type, and instance ID.

Filter tasks by status

On the **Tasks** page, the respective number of tasks in all statuses is displayed at the top. Click a task status to screen out corresponding tasks in the list. By default, tasks in the **Running** status are displayed.

Filter tasks by time

Select a time period (both the date and time can be set) in the upper-left corner of the task list to view the tasks in the corresponding time period.

Other operations

You can filter tasks, sort tasks by column, and customize columns on the Tasks page. For more information, see [Common operations](#).

3.1.5.3.4.7 DataWorks configuration

The Configuration page under DataWorks allows you to change the values of configuration items for various service roles in DataWorks.

- 1. On the Services page, click DataWorks in the left-side navigation pane.**
- 2. Select a cluster, and then click the Configuration tab.**
- 3. On the left side of the Configuration page, select a DataWorks service. The corresponding configuration items of the service are displayed on the right side.**

3.1.5.3.4.8 DataWorks cluster scaling

Apsara Bigdata Manager (ABM) supports DataWorks cluster scaling. To scale out a DataWorks cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the DataWorks cluster. To scale in a DataWorks cluster, remove physical hosts from the DataWorks cluster to the default cluster of Apsara Infrastructure Management Framework.

Background

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to the default cluster of Apsara Infrastructure Management Framework. The default cluster can be considered as an available resource pool that provides resources for scaling out business clusters. To scale out a cluster, add physical hosts in the default cluster to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster.

In the ABM console, when you scale out a DataWorks cluster, the system adds physical hosts in the default cluster to the DataWorks cluster. When you scale in a DataWorks cluster, the system removes physical hosts from the DataWorks cluster to the default cluster. The service roles of physical hosts in DataWorks include

BaseBizCdpGatewayWithNc# and BaseBizGatewayWithNc#. DataWorks cluster scaling only supports these two service roles.

Prerequisites

- **Scale-out:** The physical host to be added to a DataWorks cluster must be in the default cluster of Apsara Infrastructure Management Framework.
- **Scale-out:** If you use a host as a template host for scale-out, the service role of the host must be BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc#.
- **Scale-in:** If you use a host as a template host for scale-in, the service role of the host must be BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc#.



Note:

You can go to the MaxCompute page. Click O&M in the upper-right corner, and then click the Services tab. Click DataWorks in the left-side navigation pane, and then click the Instances tab. In the service role list, find the service role BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc#, and then click the service role name to go to the Apsara Infrastructure Management Framework console to view the hosts with the service role BaseBizCdpGatewayWithNc# or BaseBizCdpGatewayWithNc#.

Scale out a DataWorks cluster

You can add multiple hosts to a DataWorks cluster at a time to scale out the cluster. To achieve this, you need to specify an existing host as the template host. When you scale out the DataWorks cluster, the hosts copy configurations from the template host so that the hosts can be added to the cluster at a time.

1. **Go to the MaxCompute page. Click O&M in the upper-right corner, and then click the Services tab. Click DataWorks in the left-side navigation pane. Click the Slots tab, and then click the Host tab. Select a physical host with the service role BaseBizCdpGatewayWithNc# or BaseBizGatewayWithNc# as the template host.**

2. Click **Actions** in the upper-left corner, and then click **Add Server to Scale out Cluster**. In the **Add Server to Scale out Cluster** dialog box that appears, set relevant parameters.

The parameters are described as follows:

- **Refer Hostname:** the name of the template host. By default, the name of the selected host is used.
- **Hostname:** the name of the host to be added to the DataWorks cluster. Enter the name of an available host in the default cluster for scale-out. To enter multiple hostnames, separate them with commas (,).

3. Click **Run**. A message appears, indicating that the action has been submitted.
4. View the scale-out status.

Click **Actions** in the upper-left corner, and then click  next to **Add Server to Scale out Cluster** to view the scale-out history.

It may take some time for the cluster to be scaled out. In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **FAILED** indicates that the execution fails, and **SUCCESS** indicates that the execution is successful.

5. Click **Details** when the status is **RUNNING** to view the steps and progress for the execution.
6. Click **Details** when the status is **FAILED** to locate the failure cause for the failed execution. For more information, see [Locate the failure cause](#).

Scale in a DataWorks cluster

You can remove physical hosts from a DataWorks cluster to the default cluster of Apsara Infrastructure Management Framework to scale in the DataWorks cluster.

1. Go to the **MaxCompute** page. Click **O&M** in the upper-right corner, and then click the **Services** tab. Click **DataWorks** in the left-side navigation pane. Click the **Slots** tab, and then click the **Host** tab. Select one or more physical hosts with the service role **BaseBizCdpGatewayWithNc#** or **BaseBizGatewayWithNc#**.

2. Click **Actions** in the upper-left corner, and then click **Remove Server to Scale in Cluster**. In the dialog box that appears, set relevant parameters.

The parameters are described as follows:

- **Hostname:** the name of the host to be removed from the DataWorks cluster. By default, the name of the selected host is used.
- **Biz Name:** the service role of the host to be removed from the DataWorks cluster. Select the actual service role from the drop-down list. Valid values: `base-biz-cdpgatewaywithnc#` and `base-biz-gatewaywithnc#`.

3. Click **Run**. A message appears, indicating that the action has been submitted.
4. View the scale-in status.

Click **Actions** in the upper-left corner, and then click  next to **Remove Server to Scale in Cluster** to view the scale-in history.

It may take some time for the cluster to be scaled in. In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **FAILED** indicates that the execution fails, and **SUCCESS** indicates that the execution is successful.

5. Click **Details** when the status is **RUNNING** to view the steps and progress for the execution.
6. Click **Details** when the status is **FAILED** to locate the failure cause for the failed execution. For more information, see [Locate the failure cause](#).

Locate the failure cause

This section uses cluster scale-in as an example to describe how to locate the failure cause.

1. Click **Actions** in the upper-left corner, and then click  next to **Add Server to Scale out Cluster** to view the scale-out history.
2. Click **Details** for a failed execution to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

3.1.5.3.5 Tunnel service

3.1.5.3.5.1 Tunnel service O&M overview

This topic describes the concept and O&M features of the tunnel service, and how to access the tunnel service O&M page.

What is the tunnel service?

The tunnel service serves as the data tunnel of MaxCompute. You can use this service to upload data to or download data from MaxCompute.

Modules

- **Overview page:** displays the information about the tunnel service, including service overview and service status. On this page, you can also view the throughput trend chart.
- **Instances page:** displays the information about the service roles of the tunnel service.
- **Restart Tunnel Server action:** allows you to restart one or more tunnel servers.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Tunnel Service in the left-side navigation pane. The Overview page for the tunnel service appears.

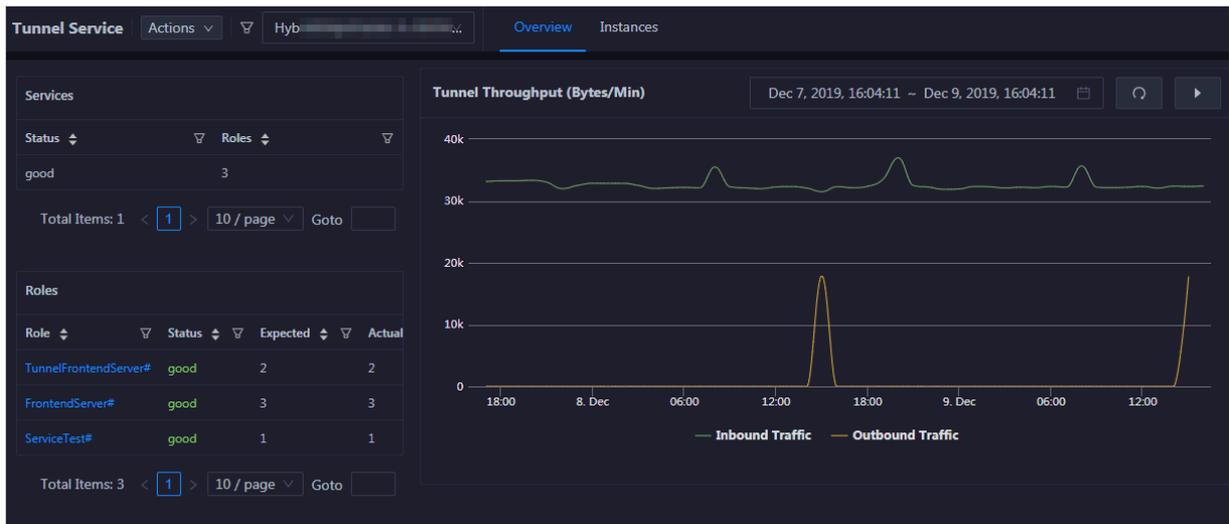


3.1.5.3.5.2 Tunnel service overview

The Overview page for the tunnel service displays key operation metrics of the tunnel service, including service overview, service status, and throughput.

Entry

On the Services page, click Tunnel Service in the left-side navigation pane. The Overview page for the tunnel service appears.



The Overview page displays key operation metrics of the tunnel service, including service overview and service status. On this page, you can also view the throughput trend chart.

Services

This section displays the numbers of available services, unavailable services, and services that are being upgraded respectively.

Roles

This section displays all service roles of the tunnel service and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

Tunnel Throughput

This chart displays the trend lines of the inbound traffic and outbound traffic in the tunnel service by minute over time in different colors.

3.1.5.3.5.3 Tunnel service instances

The Instances page displays the information about the tunnel service roles, including the name, host, IP address, status, and host status.

Entry

On the Services page, click Tunnel Service in the left-side navigation pane. Then, click the Instances tab on the right. The Instances page for the tunnel service appears.

The Instances page displays the information about all tunnel service roles, including the name, host, IP address, status, and host status. The host statuses include good, bad, and upgrading.

Other operations

You can filter or sort service roles based on a column to facilitate information retrieval on the Instances page. For more information, see [Common operations](#).

3.1.5.3.5.4 Restart tunnel servers

Apsara Bigdata Manager (ABM) allows you to restart tunnel servers for the corresponding service roles of the tunnel service.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on MaxCompute.

Context

You can restart one or more tunnel servers at a time on the Instances tab of the Tunnel Service page.

Step 1: Restart tunnel servers

1. [Log on to the ABM console](#).
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the MaxCompute page that appears, click O&M in the upper-right corner, and then click the Services tab.
4. On the Services page, click Tunnel Service in the left-side navigation pane, and then click the Instances tab.

5. On the Instances page, select one or more service roles for which you want to restart the tunnel servers, and then choose Actions > Restart Tunnel Server in the upper-left corner.
6. In the Restart Tunnel Server dialog box that appears, set relevant parameters.

The following table describes the required parameters.

Parameter	Description
Force Restart	<p>Specifies whether to forcibly restart the tunnel server for the selected service role. Valid values:</p> <ul style="list-style-type: none"> · no_force: does not forcibly restart the tunnel server. If a service role is in the running state, the corresponding tunnel server is not restarted. · force: forcibly restarts the tunnel server. A tunnel server is restarted no matter which state the corresponding service role is in.
Hostname	<p>The hostname of the selected service role. Multiple hostnames are separated with commas (.). The value is automatically filled. You cannot specify a value for this parameter.</p>

7. Click Run. A message appears, indicating that the action has been submitted.

Step 2: View the execution status or progress

1. On the Overview tab or the Instances tab of the Tunnel Service page, click Actions in the upper-left corner. Then, click  next to Restart Tunnel Server to view the execution history.

In the Current Status column, **RUNNING** indicates that the execution is in progress, **FAILED** indicates that the execution fails, and **SUCCESS** indicates that the execution is successful.

2. Click Details in the Details column of a task in the **RUNNING** state to view the execution progress.

(Optional) Step 3: Locate the failure cause

If the status of the task is **FAILED**, you can view the execution logs to locate the failure cause.

1. On the Overview tab or the Instances tab of the Tunnel Service page, click **Actions** in the upper-left corner. Then, click  next to **Restart Tunnel Server** to view the execution history.
2. In the execution history dialog box, click **Details** in the **Details** column of the task to view the details.
3. On the Servers tab of the failed step, click **View Details** in the **Actions** column of a failed server. The **Execution Output** tab appears in the **Execution Details** section. You can view the output to locate the failure cause.

3.1.5.4 Cluster O&M

3.1.5.4.1 Cluster O&M overview

This topic describes the cluster O&M features of MaxCompute and how to access the cluster O&M page.

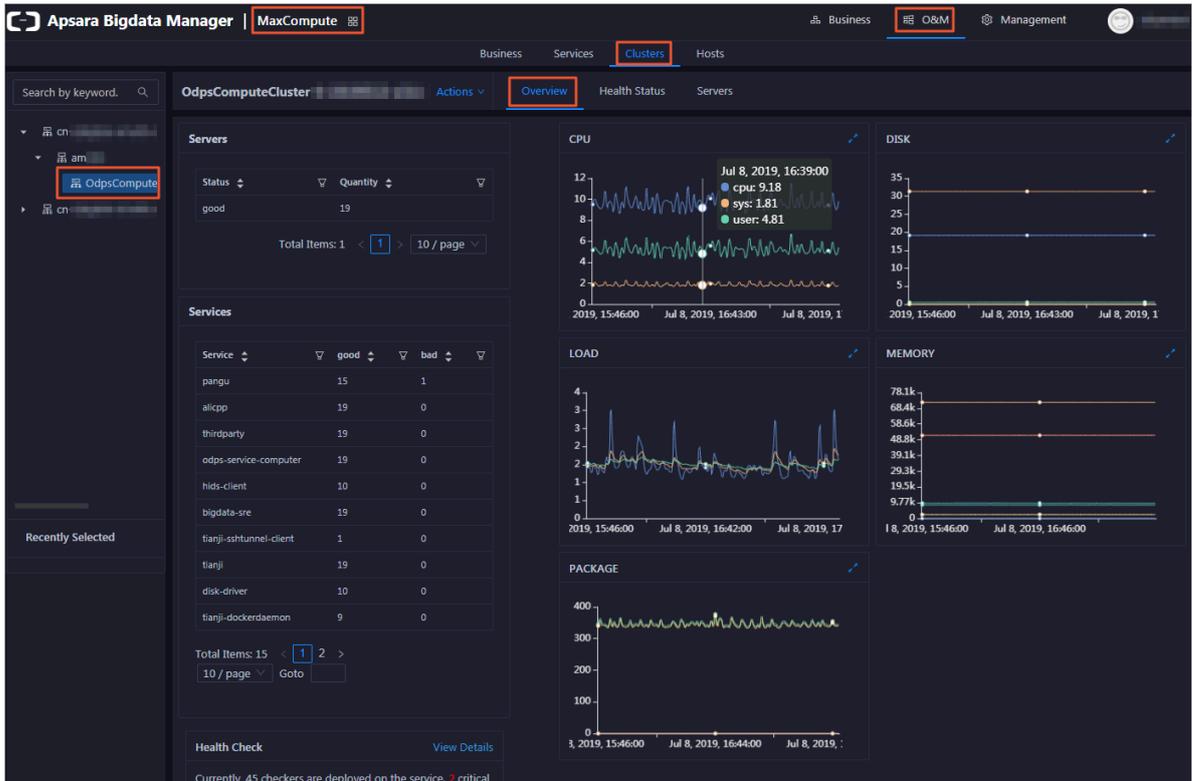
Cluster O&M features

- **Overview page:** displays the overall running information about a cluster. On this page, you can view the host status, service status, health check result, and health check history. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.
- **Health Status page:** displays the check results of all checkers in a cluster. The check results include the Critical, Warning, Exception, and OK types.
- **Servers page:** displays information about all hosts in a cluster, including the CPU usage, memory usage, root disk usage, packet loss rate, and packet error rate.
- **Cluster scaling action:** allows you to scale out or scale in a MaxCompute cluster by adding or removing physical hosts.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click **MaxCompute**.
3. On the page that appears, click **O&M** at the top, and then click the **Clusters** tab.

4. On the Clusters page, select a cluster in the left-side navigation pane. The Overview page for the cluster appears.

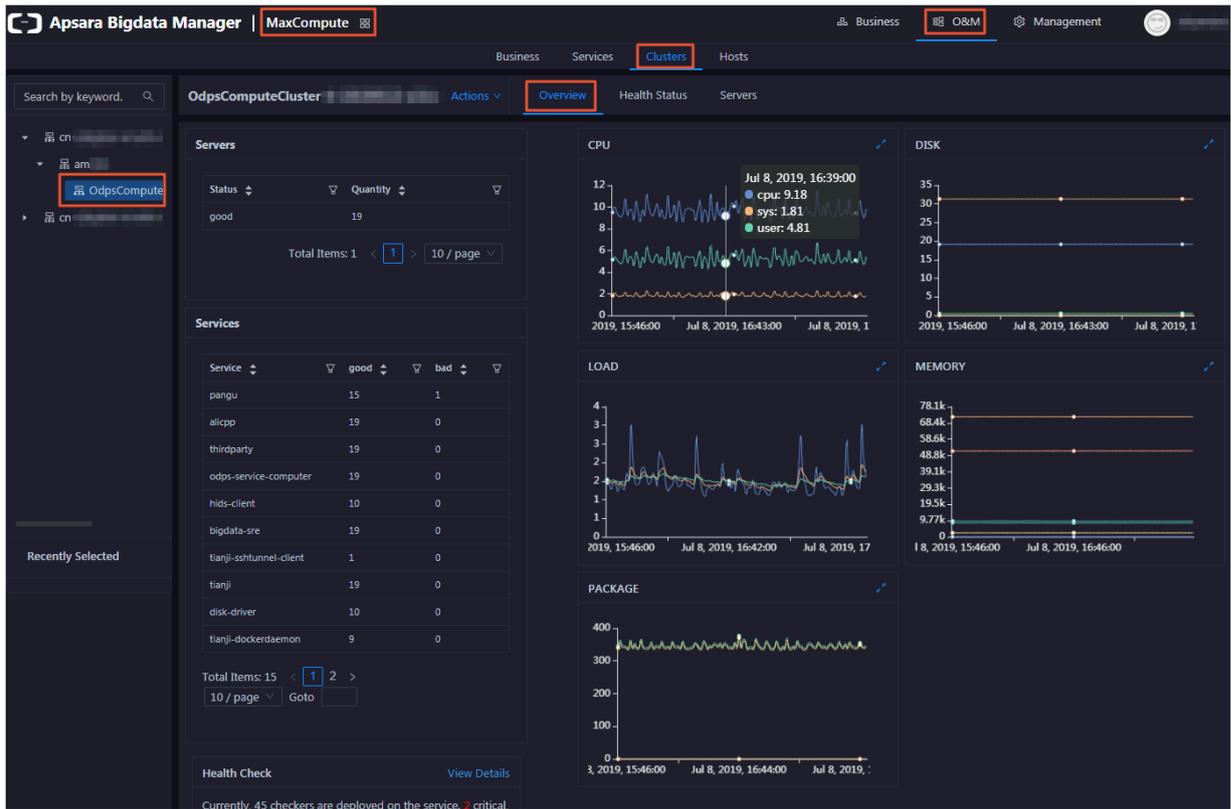


3.1.5.4.2 Cluster overview

The cluster overview page displays key metrics for a cluster. On this page, you can view the host status, service status, health check result, and health check history. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

Entry

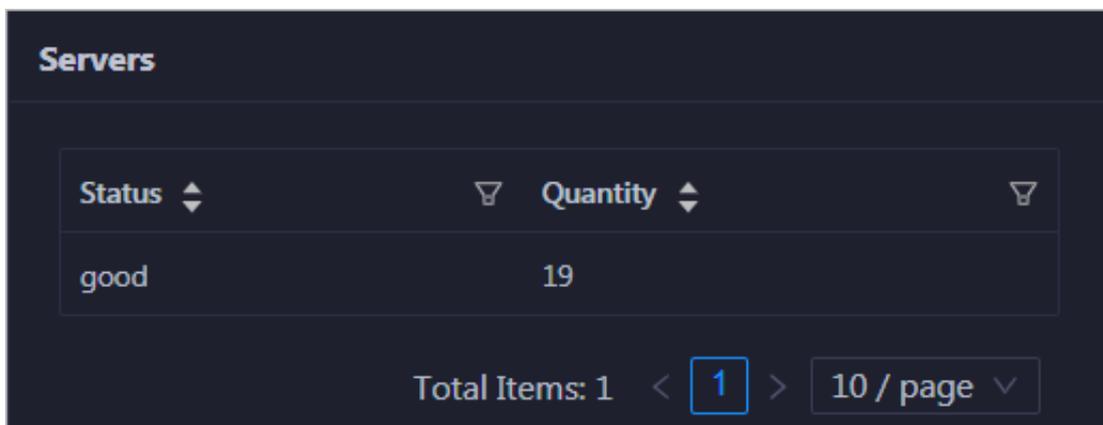
On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.



On the Overview page, you can view the host status, service status, health check result, and health check history of the selected cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage. To view information about a cluster, select a region in the left-side navigation pane, select an IDC in the region, and then select a cluster in the IDC.

Servers

This section displays all host statuses and the number of hosts in each status. The host statuses include good and bad.



Services

This section displays all services deployed in the cluster and the respective number of available and unavailable services.

Service	good	bad
pangu	15	1
alicpp	19	0
thirdparty	19	0
odps-service-computer	19	0
hids-client	10	0
bigdata-sre	19	0
tianji-sshtunnel-client	1	0
tianji	19	0
disk-driver	10	0
tianji-dockerdaemon	9	0

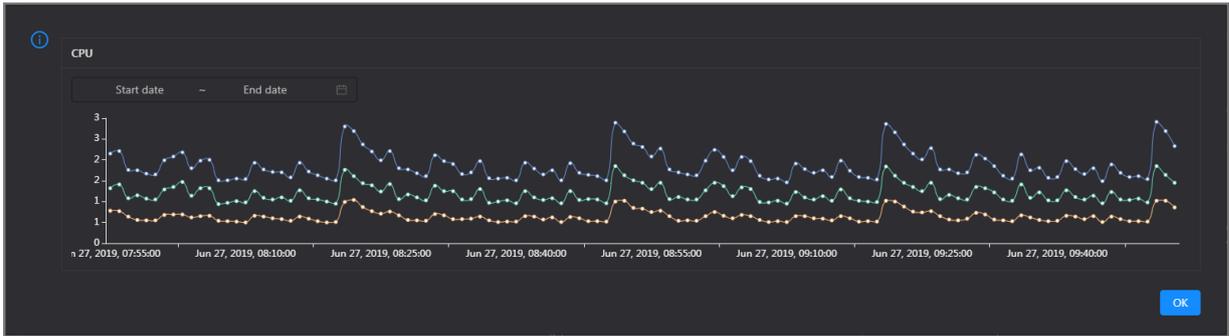
Total Items: 15 < 1 2 > 10 / page Goto

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

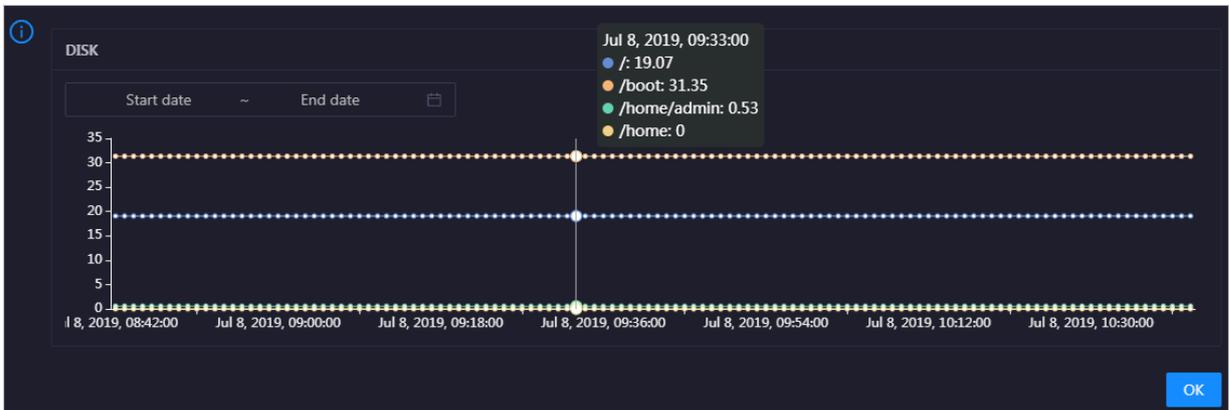
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

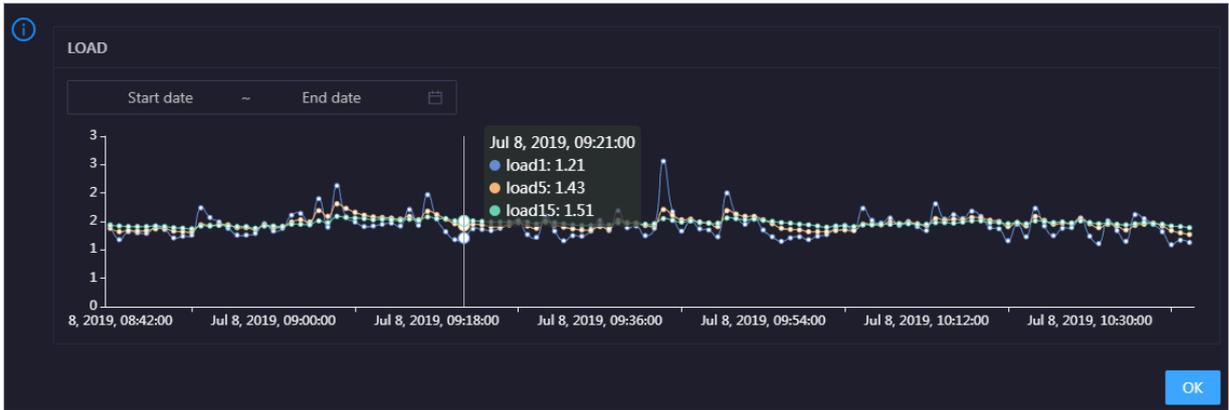


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

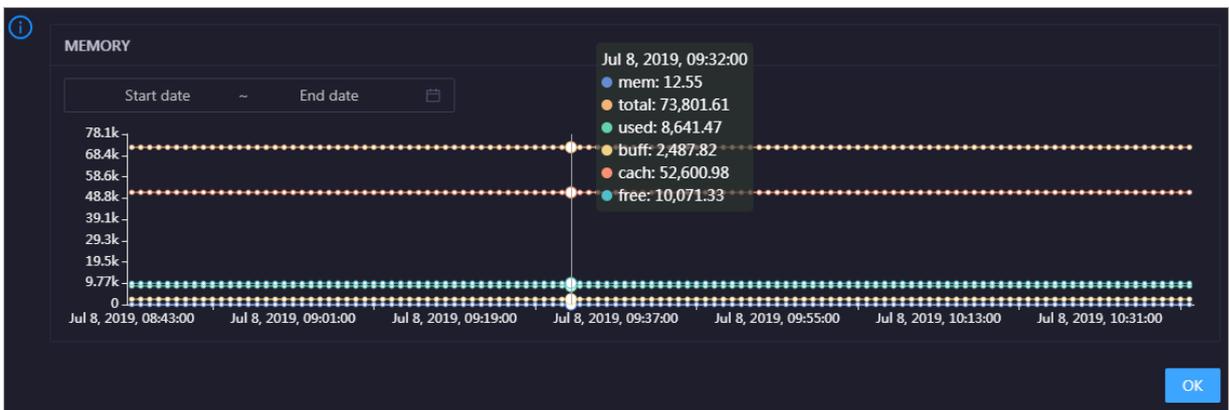


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

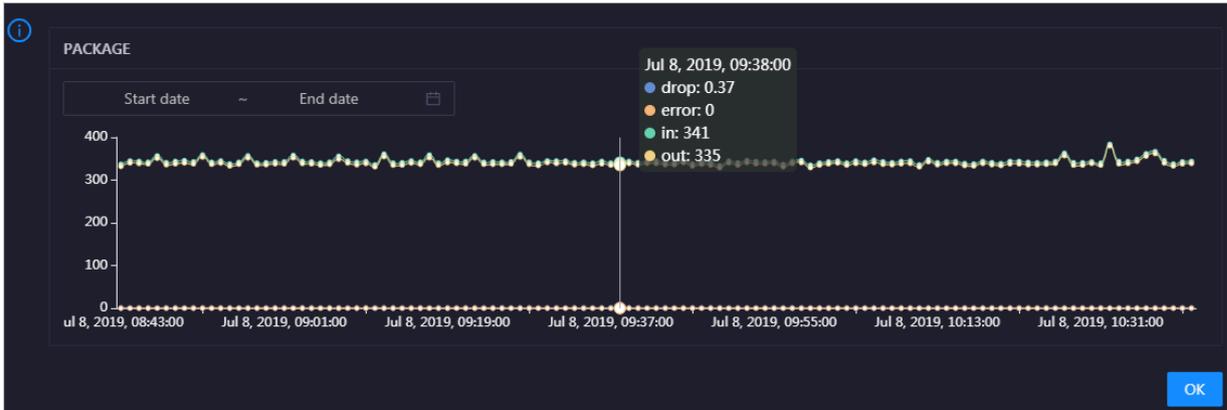


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

Health Check

This section displays the number of checkers deployed for the cluster and the respective number of hosts with Critical, Warning, and Exception alerts.

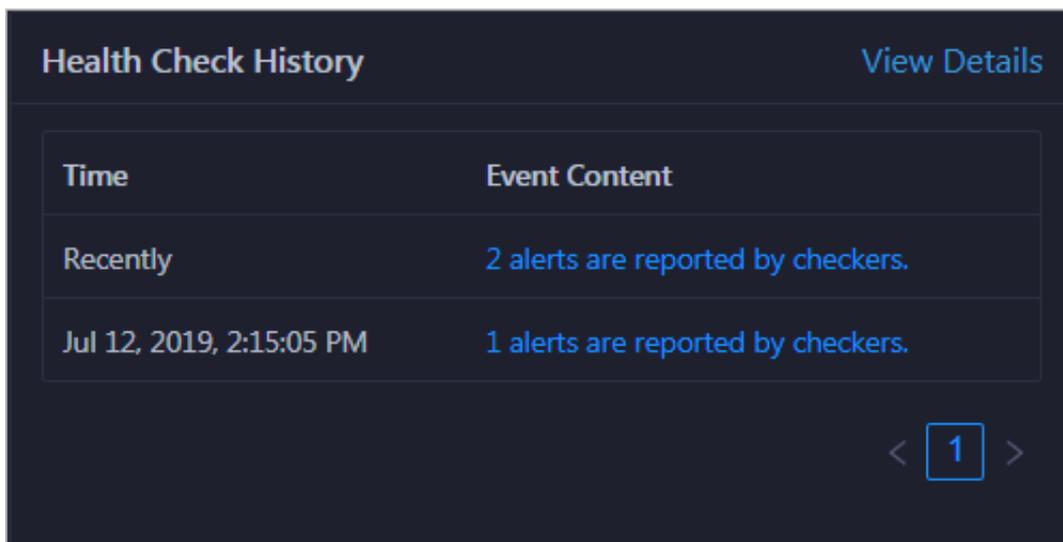
Health Check [View Details](#)

Currently, 45 checkers are deployed on the service. 2 critical, 0 exception, and 11 warning alerts are reported.

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

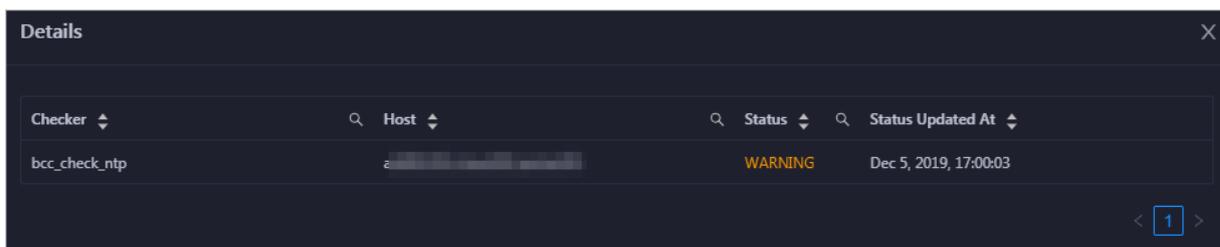
Health Check History

This section displays a record of the health checks performed on the cluster. You can view the respective number of Critical, Warning, and Exception alerts for each health check.



Click **View Details** to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

You can click the event content of a check to view the exception items.



3.1.5.4.3 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.

Checker	Source	Critical	Warning	Exception	Actions
+ eodps_check_nuwa	tcheck	1	0	0	Details
+ eodps_check_aas	tcheck	1	0	0	Details
+ bcc_check_ntp	tcheck	0	10	0	Details
+ eodps_check_schedulerpoolsize	tcheck	0	1	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_host_live_check	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details
+ bcc_network_tcp_connections_checker	tcheck	0	0	0	Details

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

Name: bcc_tsar_tcp_checker **Source:** tcheck

Alias: TCP Retransmission Check **Application:** bcc

Type: system **Scheduling:** Enable

Data Collection: Enable

Default Execution Interval: 0 0/5 * * * ?

Description:
 This checker uses tsar commands to test the retransmission rate. Reason: Server overloads or network fluctuations. Fix:

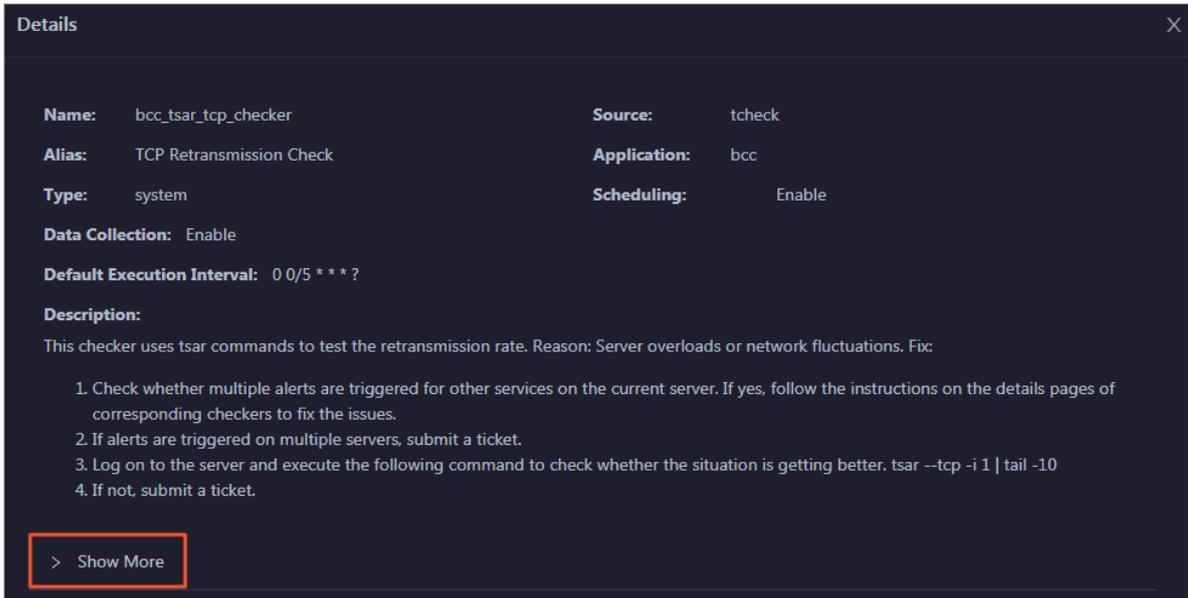
1. Check whether multiple alerts are triggered for other services on the current server. If yes, follow the instructions on the details pages of corresponding checkers to fix the issues.
2. If alerts are triggered on multiple servers, submit a ticket.
3. Log on to the server and execute the following command to check whether the situation is getting better. `tsar --tcp -i 1 | tail -10`
4. If not, submit a ticket.

> Show More

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is

enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

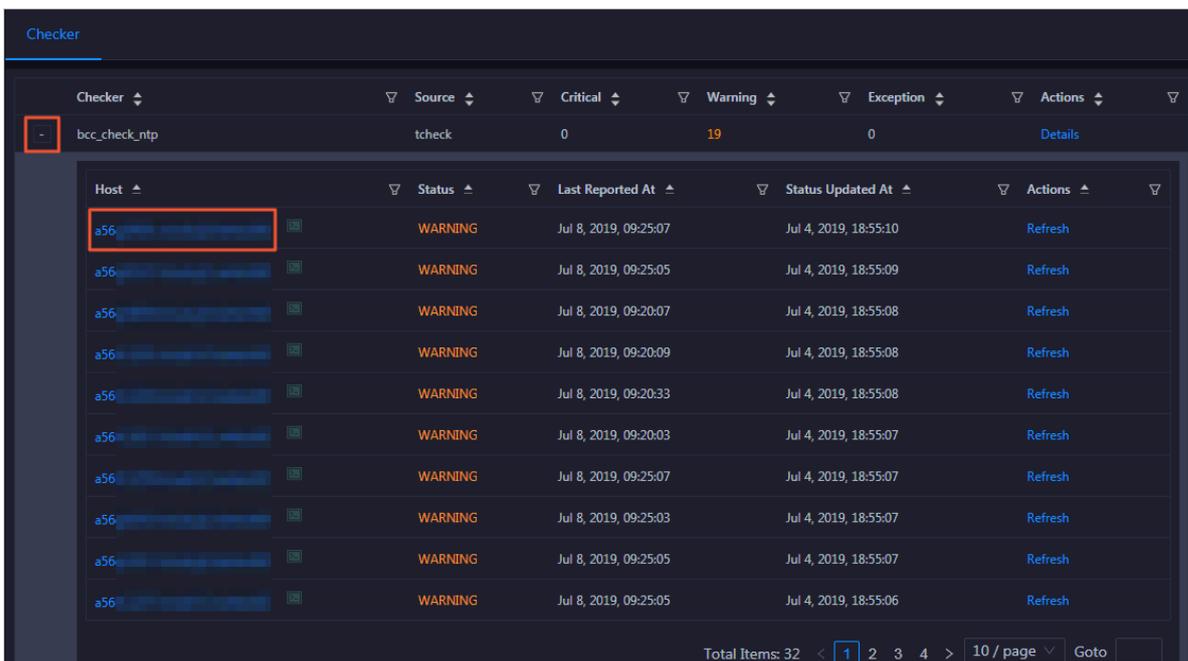


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

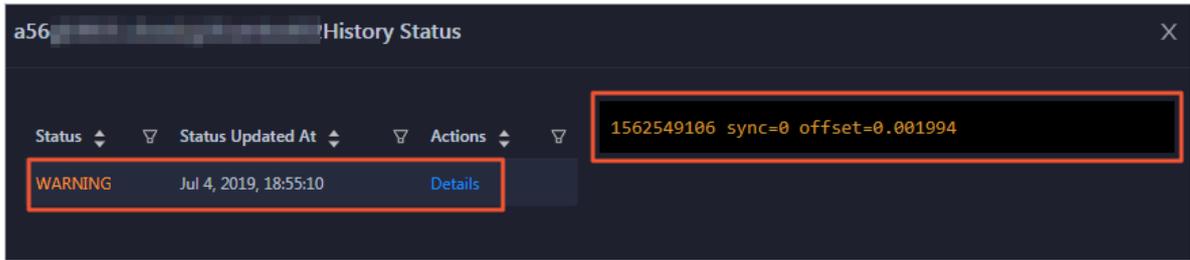
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.

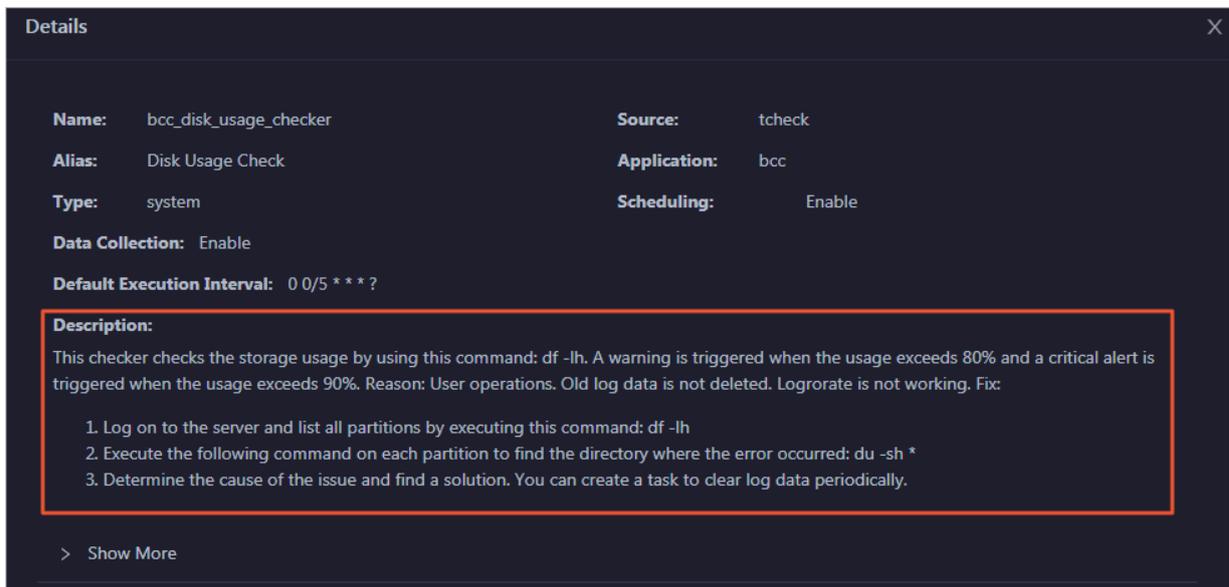


2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

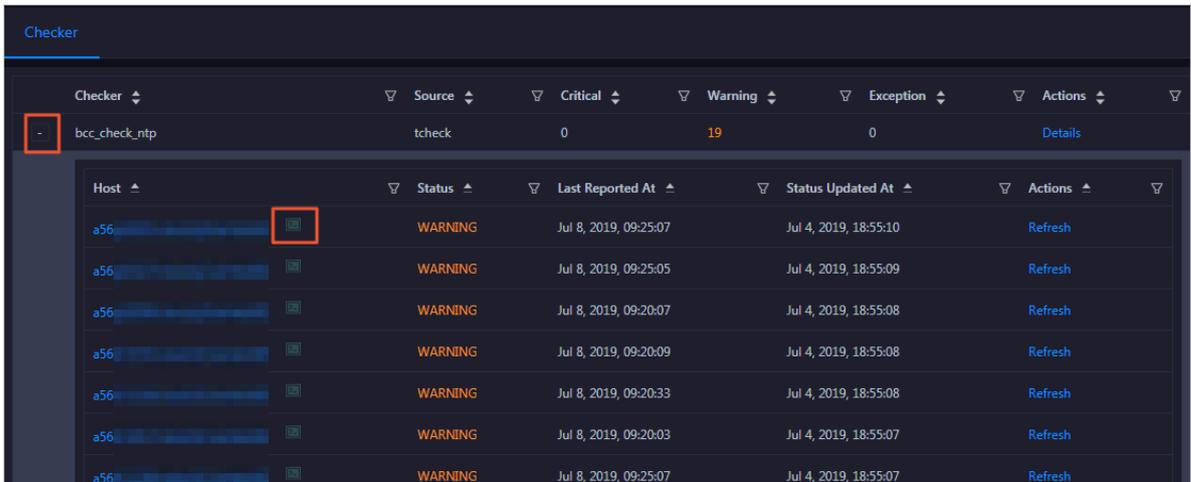
- On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



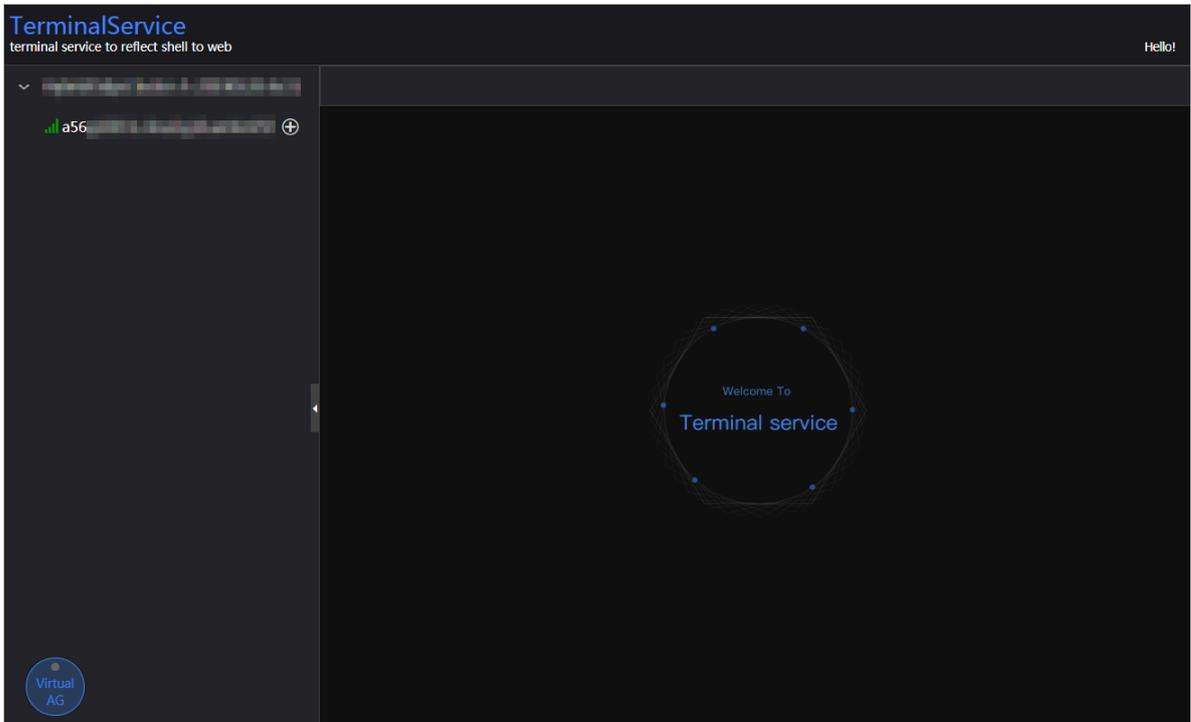
Log on to a host

- To log on to a host to clear alerts or perform other operations, follow these steps:

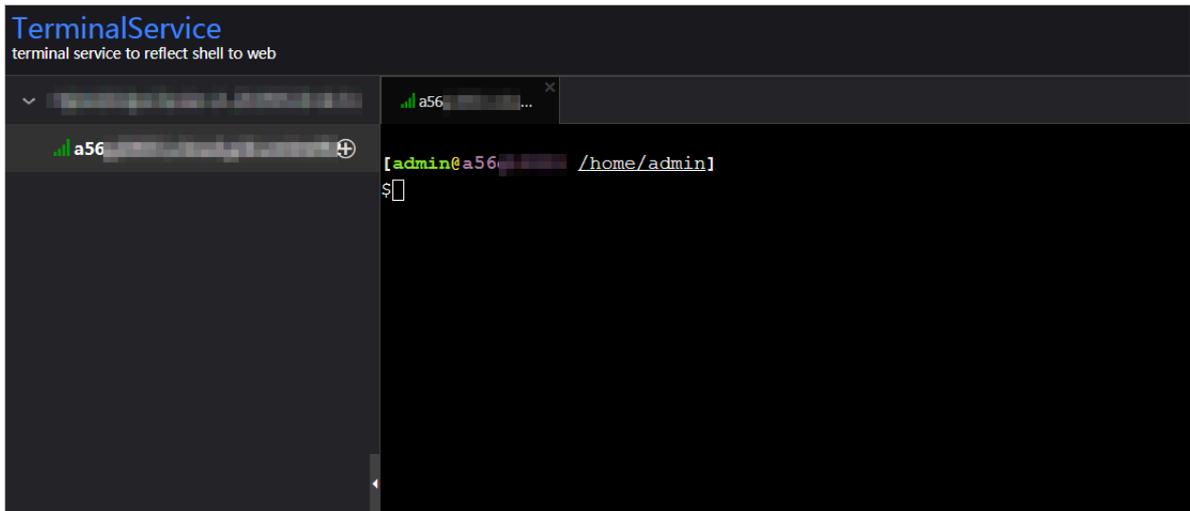
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.



3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

 The screenshot shows the "Checker" page in a web interface. At the top, there are filters for "Checker", "Source", "Critical", "Warning", "Exception", and "Actions". Below these filters is a table with the following data:

Checker	Source	Critical	Warning	Exception	Actions
bcc_check_ntp	tcheck	0	19	0	Details

 Below this is another table listing hosts:

Host	Status	Last Reported At	Status Updated At	Actions
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

 The "Refresh" button in the first row of the second table is highlighted with a red box.

3.1.5.4.4 Cluster hosts

The cluster host page displays information about hosts, including the hostname, IP address, role, type, CPU usage, total memory size, available memory size, load, root disk usage, packet loss rate, and packet error rate.

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Servers tab. The Servers page for the cluster appears.

Hostname	IP	Role	Type	CPU Usage (%)	Total Memory (MB)	Idle Memory (MB)	Load1	Root Disk Usage (%)	Packet Loss Rate	Packet Error Rate
a56	10.	BigGraphWorker	Q41.2B	1	270685.86	225428.58	0.3	24.7	0	0
a56	10.	BigGraphWorker	Q41.2B	1.1	270685.86	222629.45	0.2	24.6	0	0
a56	10.	BigGraphWorker	Q41.2B	1	270685.86	219430.3	0.2	24.6	0	0
a56	10.	OdpsComputer	Q45.2B	1.1	115866.53	13021.39	0.7	26.5	0	0
a56	10.	OdpsComputer	Q45.2B	1.2	115866.53	14423.42	0.2	26.2	0	0
a56	10.	OdpsComputer	Q45.2B	1.3	115866.53	11324.58	0.6	26.3	0	0
a56	10.	OdpsComputer	Q45.2B	1.6	115866.53	15583.15	0.5	26.2	0	0
a56	10.	OdpsComputer	Q45.2B	1.5	115866.53	8582.05	0.5	26.5	0	0
a56	10.	OdpsComputer	Q45.2B	1.5	115866.53	14608.04	1	26.4	10	0
a56	10.	OdpsComputer	Q45.2B	2	115866.53	7033.77	0.9	26.2	0	0

To view more information about a host, click the name of the host. The [Host overview](#) page appears.

3.1.5.4.5 Cluster scaling

Apsara Bigdata Manager (ABM) supports MaxCompute cluster scaling. To scale out a MaxCompute cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the MaxCompute cluster. To scale in a MaxCompute cluster, remove physical hosts from the MaxCompute cluster to the default cluster of Apsara Infrastructure Management Framework.

Background

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to the default cluster of Apsara Infrastructure Management Framework. The default cluster can be considered as an available resource pool that provides resources for scaling out business clusters. To scale out a cluster, add physical hosts in the default cluster to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster.

In the ABM console, when you scale out a MaxCompute cluster, the system adds physical hosts in the default cluster to the MaxCompute cluster. When you scale in a MaxCompute cluster, the system removes physical hosts from the MaxCompute cluster to the default cluster.

Prerequisites

- **Scale-out:** The physical host to be added to a MaxCompute cluster must be an SInstance host in the default cluster of Apsara Infrastructure Management Framework.
- **Scale-out:** If you use a host as a template host for scale-out, the host must be an SInstance host. You can log on to the admingateway host in the MaxCompute cluster to view SInstance hosts.
- **Scale-in:** The physical host to be removed from a MaxCompute cluster must be an SInstance host. You can log on to the admingateway host in the MaxCompute cluster to view SInstance hosts.

Scale out a MaxCompute cluster

You can add multiple hosts to a MaxCompute cluster at a time to scale out the cluster. To achieve this, you need to specify an existing host as the template host. When you scale out the MaxCompute cluster, the hosts copy configurations from the template host so that the hosts can be added to the cluster at a time.

1. Log on to the admingateway host in the MaxCompute cluster. Run the `r ttrtl` command to view SInstance hosts. For more information about how to log on to a host, see [Log on to a host](#).

```

TerminalService
terminal service to reflect shell to web

[admin@vm /home/admin]
$ r ttrtl
total tubo in cluster=11

detail table for every machine:
Machine Name | CPU | Memory | Other
a56f11 | 3,900 | 235,048 | BigGraphInstance:99
a56f11 | 3,900 | 235,048 | BigGraphInstance:99
a56e09 | 3,900 | 167,510 | OdpsSpecialInstance:20 OdpsCommonInstance:20
a56e09 | 3,900 | 235,048 | BigGraphInstance:99
a56f11 | 3,900 | 167,510 | GraphInstance:8 RTInstance:4 SInstance:99
a56f11 | 3,900 | 167,510 | GraphInstance:8 RTInstance:4 SInstance:99
a56e09 | 3,900 | 167,510 | GraphInstance:8 RTInstance:4 SInstance:99
a56e09 | 3,900 | 167,510 | OdpsSpecialInstance:20 OdpsCommonInstance:20
a56e09 | 3,900 | 167,510 | GraphInstance:8 RTInstance:4 SInstance:99
a56e07 | 3,900 | 167,510 | GraphInstance:8 RTInstance:4 SInstance:99
a56f11 | 3,900 | 167,510 | GraphInstance:8 RTInstance:4 SInstance:99
Total | 42,900 | 2,045,224 | NA

[admin@vm /home/admin]
$
    
```

2. On the Clusters page, select a cluster in the left-side navigation pane, and click the Servers tab. On the page that appears, select an SInstance host as the template host.

Hostname	IP	Role	Type	CPU Usage (%)	Total Memory (MB)	Idle Memory (MB)	Load1	Root Disk Usage (%)	Packet Loss Rate	Packet Error Rate
a5-...	10...	OdpsComputer	Q45.2B	1.1	115866.53	14561.63	0.6	26.4	11	0
a5-...	10...	OdpsComputer	Q45.2B	0.9	115866.53	13007.87	0.4	26.5	0	0
a5-...	10...	OdpsComputer	Q45.2B	1.1	115866.53	14446.09	0.2	26.2	0	0
a5-...	10...	OdpsComputer	Q45.2B	1.2	115866.53	15602.31	0.8	26.2	0	0
a5-...	10...	OdpsComputer	Q45.2B	1.5	115866.53	7069.95	0.6	26.2	0	0
a5-...	10...	OdpsController	Q45.2B	4.3	115866.53	4605.41	3	34.1	0	0
a5-...	10...	OdpsController	Q45.2B	2	115866.53	4515.82	1.2	34.4	0	0
a5-...	10...	TunnelFrontendServer	Q45.2B	1.4	115866.53	7414.54	0.7	26.8	0	0
a5-...	10...	TunnelFrontendServer	Q45.2B	1.7	115866.53	10613.69	0.8	27	0	0
vn-...	10...	PanguMaster	VM	11.4	54108	238.52	1.6	11.7	0	0

Total Items: 31 < 1 2 3 4 > 10 / page Goto

3. Click Actions in the upper-left corner, and then click Scale out Cluster. In the Scale out Cluster dialog box that appears, set relevant parameters.

The parameters are described as follows:

- **Refer_hostname:** the name of the template host. By default, the name of the selected host is used.
- **Hostname:** the name of the host to be added to the MaxCompute cluster. The drop-down list displays all available hosts in the default cluster for scale-out. You can select one or more hosts from the drop-down list.

4. Click OK. A message appears, indicating that the action has been submitted.
5. View the scale-out status.

Click Actions in the upper-left corner, and then click Execution History next to Scale out Cluster to view the scale-out history.

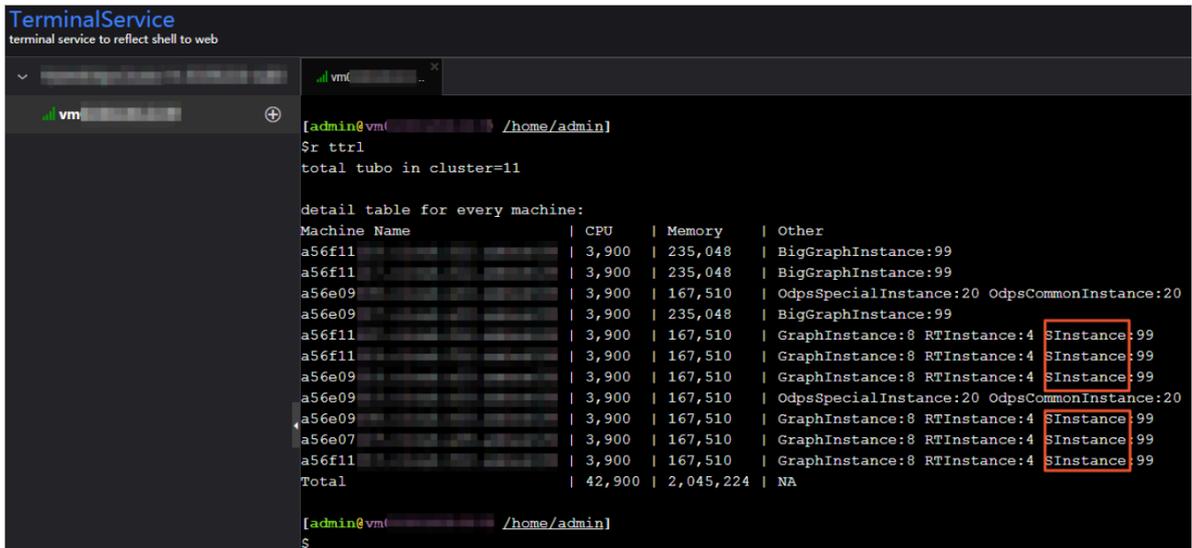
It may take some time for the cluster to be scaled out. In the Current Status column, RUNNING indicates that the execution is in progress, FAILED indicates that the execution fails, and SUCCESS indicates that the execution is successful.

6. Click Details when the status is RUNNING to view the steps and progress for the execution.
7. Click Details when the status is FAILED to locate the failure cause for the failed execution. For more information, see [Locate the failure cause](#).

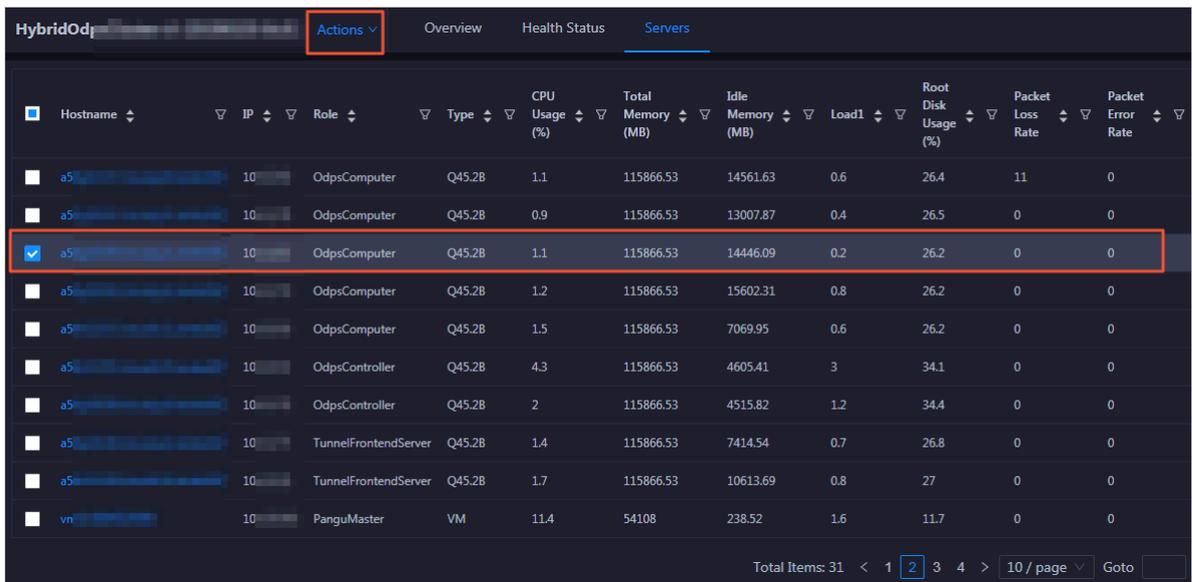
Scale in a MaxCompute cluster

You can remove multiple hosts from a MaxCompute cluster at a time to scale in the cluster.

1. Log on to the admingateway host in the MaxCompute cluster. Run the `rsctl` command to view SInstance hosts. For more information about how to log on to a host, see [Log on to a host](#).



2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Servers tab. On the page that appears, select one or more SInstance hosts to be removed.



3. Click **Actions** in the upper-left corner, and then click **Scale in Cluster**. In the **Scale in Cluster** dialog box that appears, set **Hostname**.

Hostname: the name of the host to be removed from the MaxCompute cluster. By default, the name of the selected host is used.

4. Click **Run**. A message appears, indicating that the action has been submitted.
5. View the scale-in status.

Click **Actions** in the upper-left corner, and then click **Execution History** next to **Scale in Cluster** to view the scale-in history.

It may take some time for the cluster to be scaled in. In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **FAILED** indicates that the execution fails, and **SUCCESS** indicates that the execution is successful.

6. Click **Details** when the status is **RUNNING** to view the steps and progress for the execution.
7. Click **Details** when the status is **FAILED** to locate the failure cause for the failed execution. For more information, see [Locate the failure cause](#).

Locate the failure cause

This section uses cluster scale-in as an example to describe how to locate the failure cause.

1. On the **Clusters** page, click **Actions** in the upper-left corner, and then click **Execution History** next to **Scale in Cluster** to view the scale-in history.
2. Click **Details** for a failed execution to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

3.1.5.5 Host O&M

3.1.5.5.1 Host O&M overview

This topic describes the host O&M features of MaxCompute and how to access the host O&M page.

Host O&M features

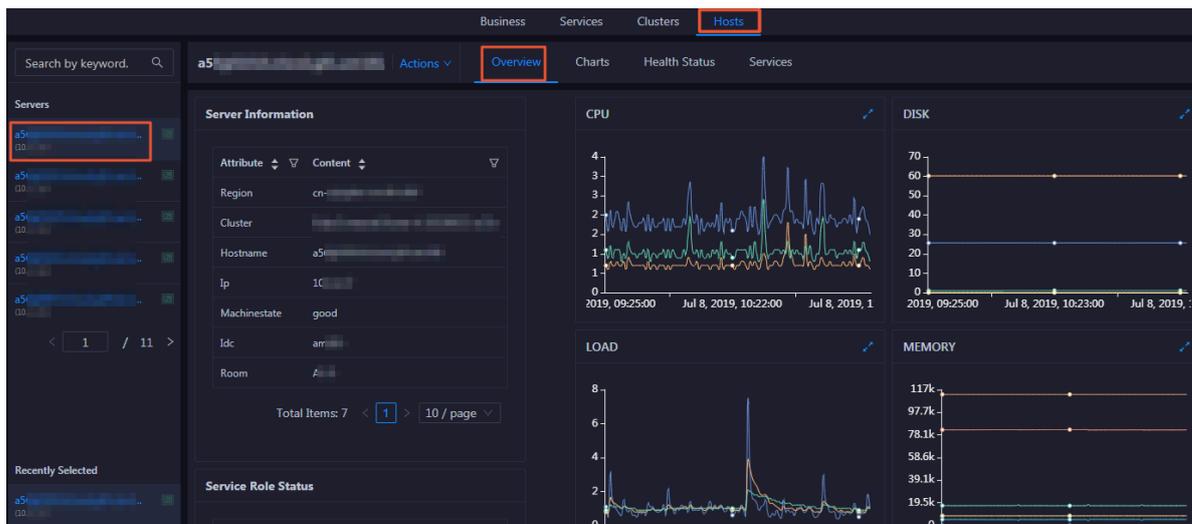
- **Overview page:** displays brief information about a host in a MaxCompute cluster . On this page, you can view the attributes, services, service roles, health check

result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

- **Charts page:** displays the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.
- **Health Status page:** displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.
- **Services page:** displays the cluster, service instances, and service instance roles of a host.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click MaxCompute.
3. On the page that appears, click O&M at the top, and then click the Hosts tab.
4. On the Hosts page, select a host in the left-side navigation pane. The Overview for the host appears.

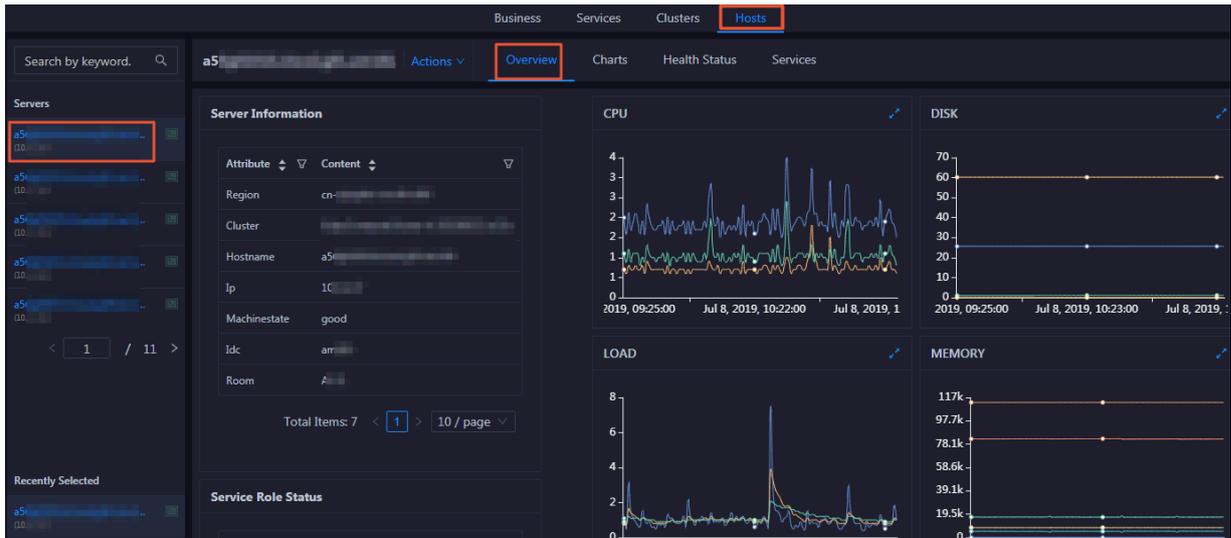


3.1.5.5.2 Host overview

The host overview page displays brief information about a host in a MaxCompute cluster. On this page, you can view the attributes, services, service roles, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

Entry

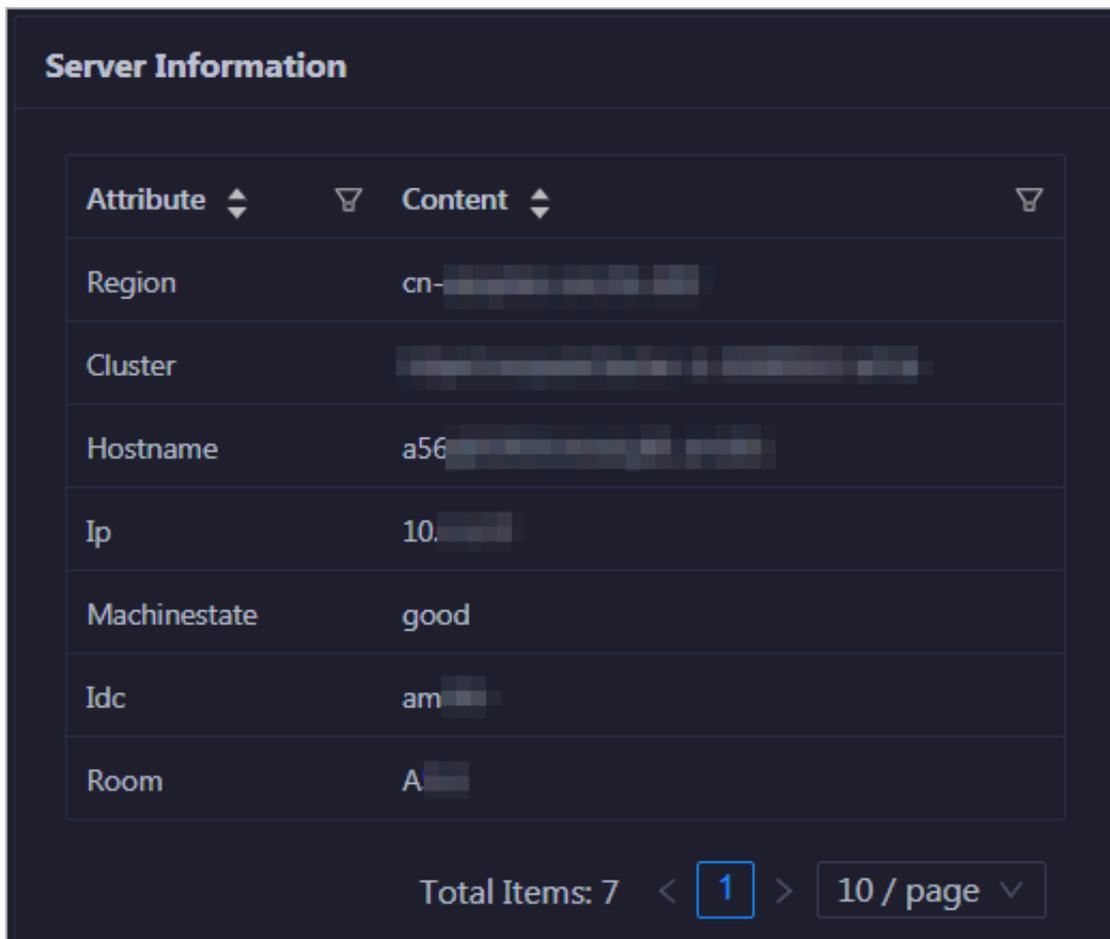
On the Hosts page, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.



On the Overview page, you can view the server information, service roles, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

Server Information

This section displays the information about the host, including the region, cluster, name, IP address, status, IDC, and server room of the host.



Service Role Status

This section displays the information about the services deployed on the host, including the roles, statuses, and number of services.

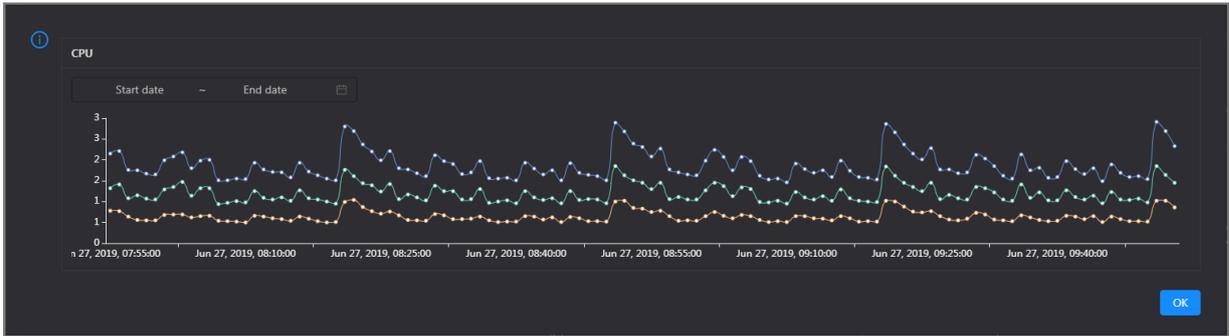
Service	Role	State	Num
alicpp	OdpsRpm#	good	1
bigdata-sre	Agent#	good	1
disk-driver	DiskDriverWorker#	good	1
hids-client	HidsClient#	good	1
nuwa	NuwaConfig#	good	1
odps-service-computer	PackageInit#	good	1
odps-service-frontend	TunnelFrontendServer#	good	1
thirdparty	ThirdpartyLib#	good	1
tianji	TianjiClient#	good	1
pangu	PanguChunkserver#	good	1

Total Items: 19 < 1 2 > 10 / page Goto

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

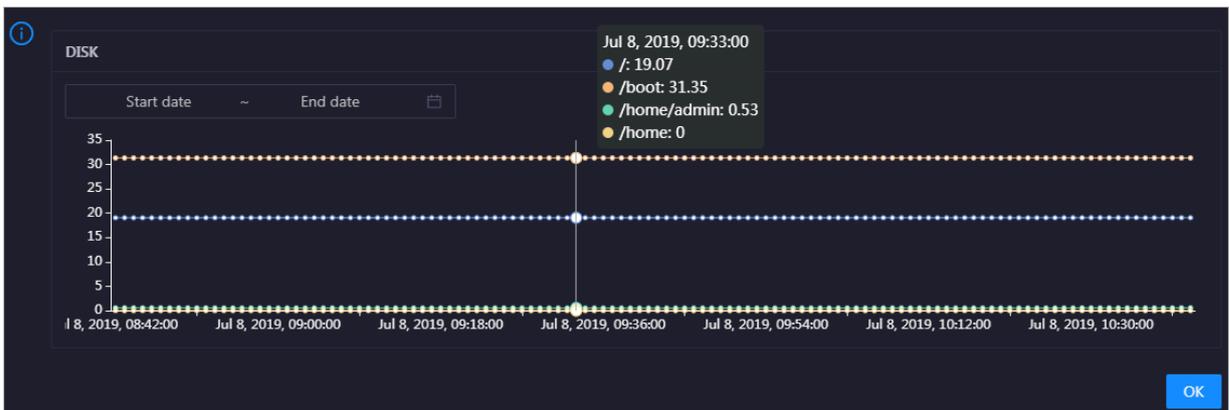


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

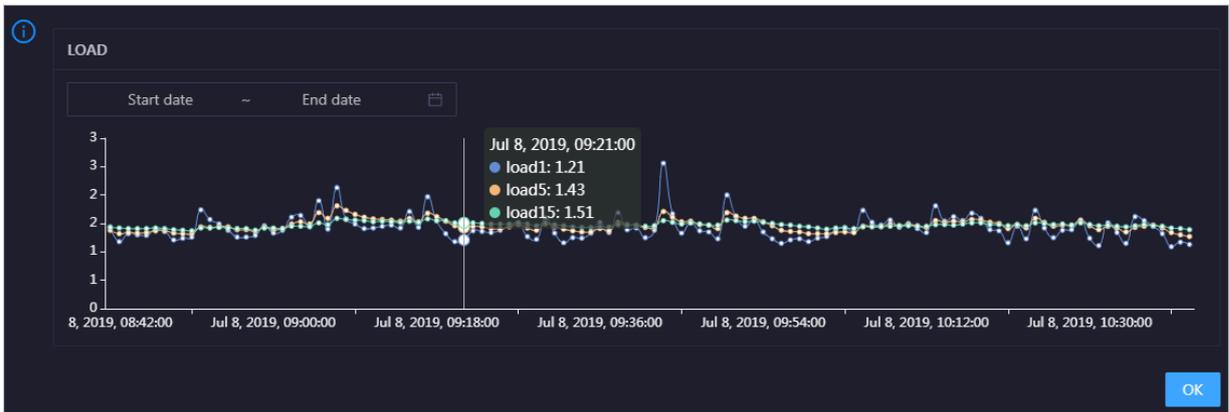


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

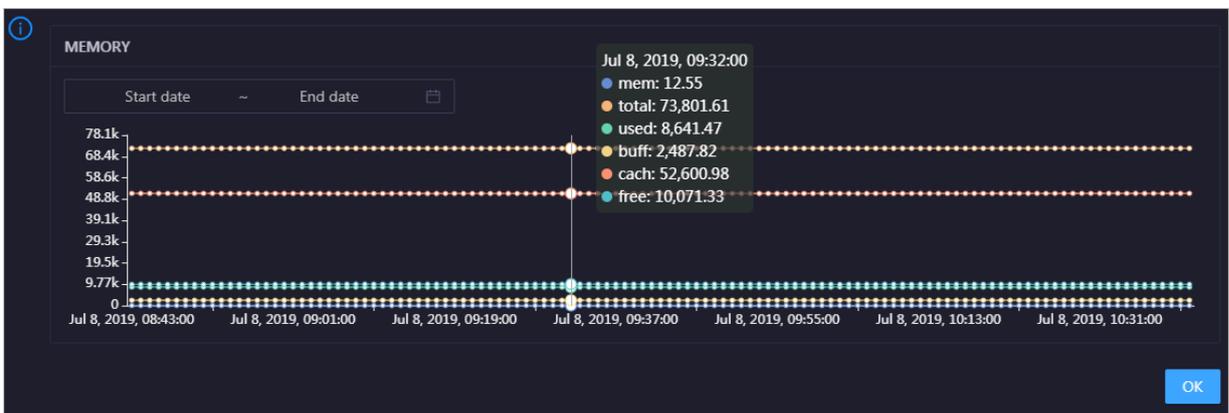


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

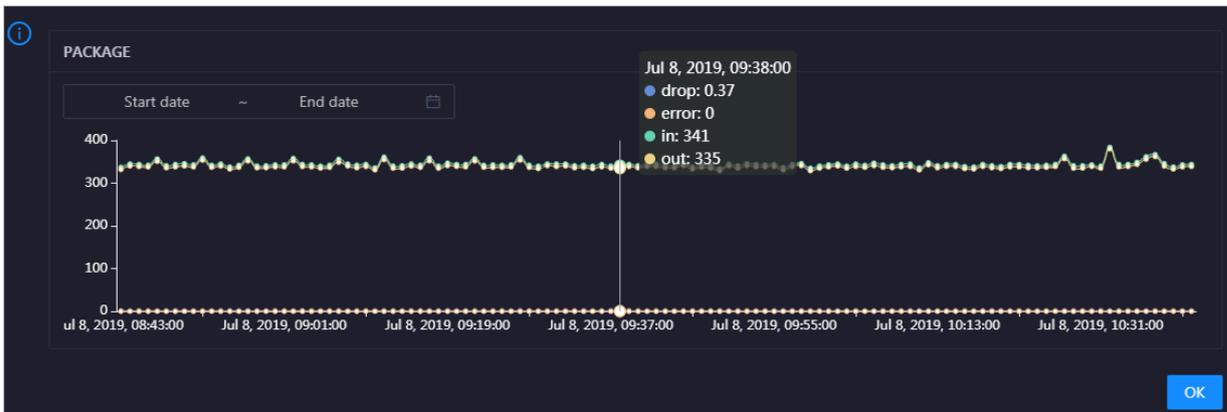


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

Health Check [View Details](#)

Currently, 10 checkers are deployed on the service. 0 critical, 0 exception, and 1 warning alerts are reported.

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

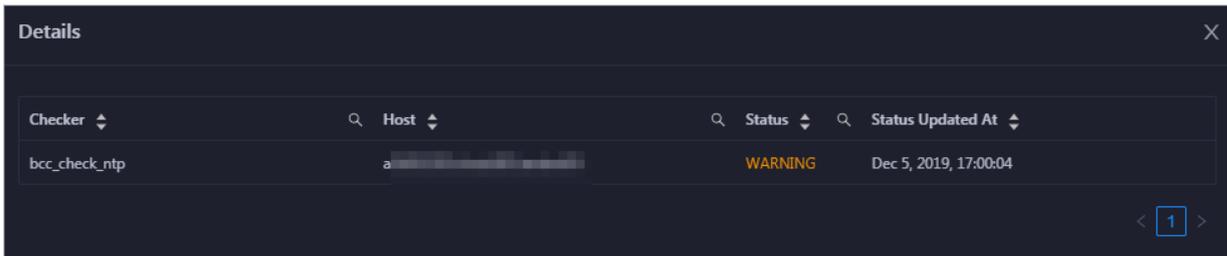
Health Check History

This section displays a record of the health checks performed on the host.

Health Check History		View Details
Time	Event Content	
Recently	1 alerts are reported by checkers.	< 1 >

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

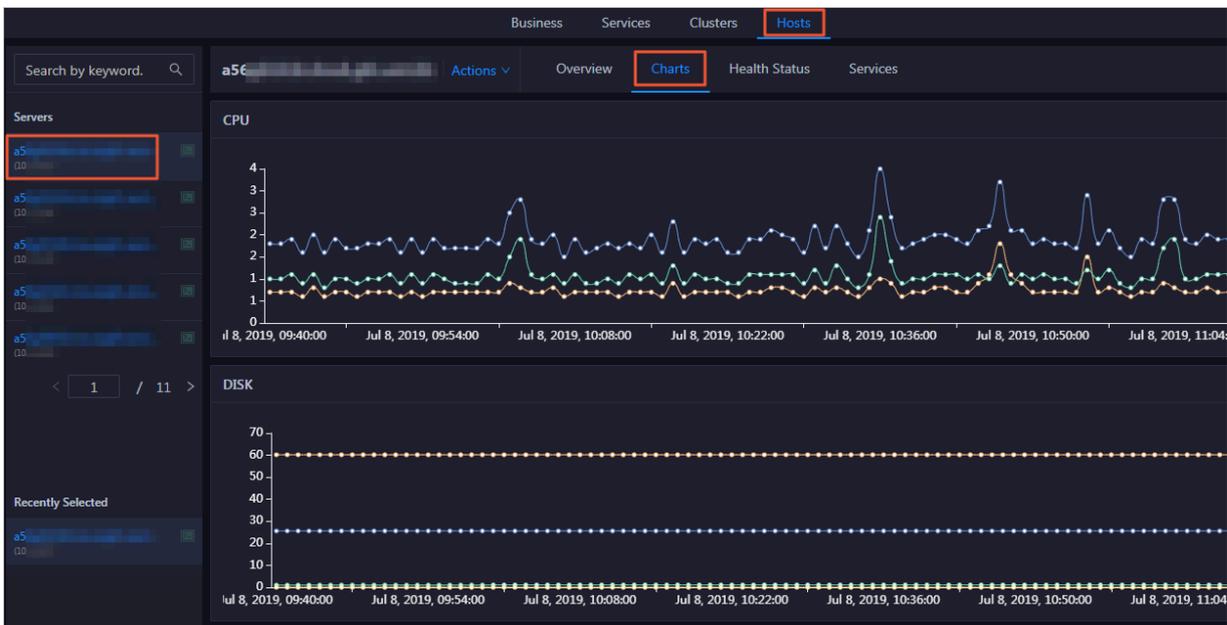
You can click the event content of a check to view the exception items.



3.1.5.5.3 Host charts

On the host chart page, you can view the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.

On the Hosts page, select a host in the left-side navigation pane, and then click the Charts tab. The Charts page for the host appears.



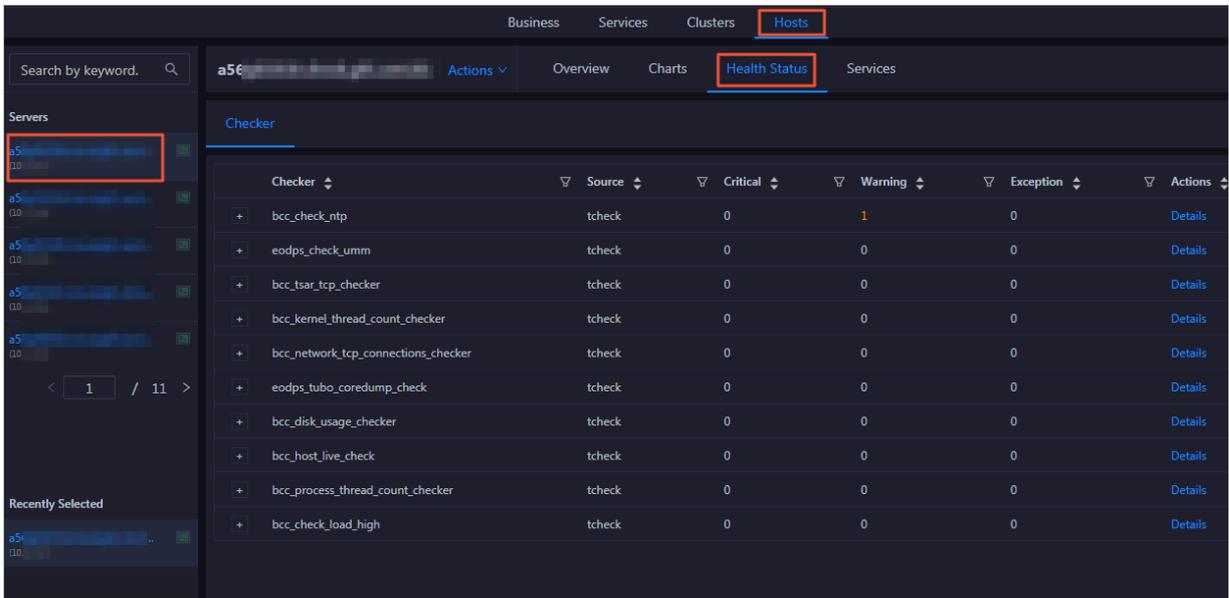
The Charts page displays trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host. For more information, see [Host overview](#).

3.1.5.5.4 Host health

On the Health Status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

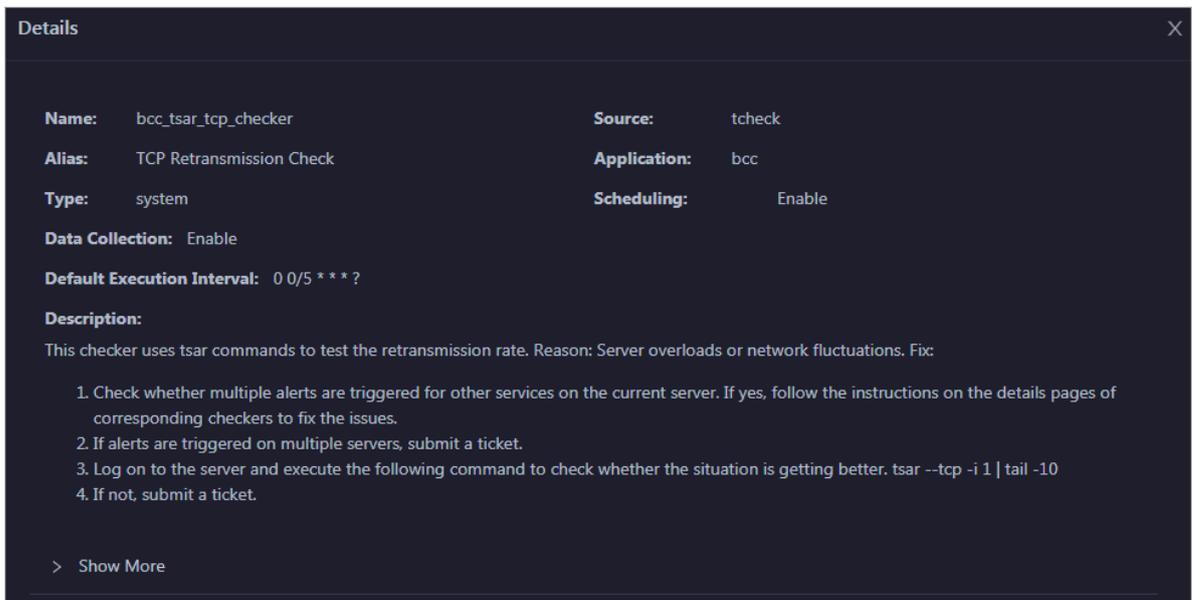
On the Hosts page, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.



On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into Critical, Warning, Exception, and OK. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

View checker details

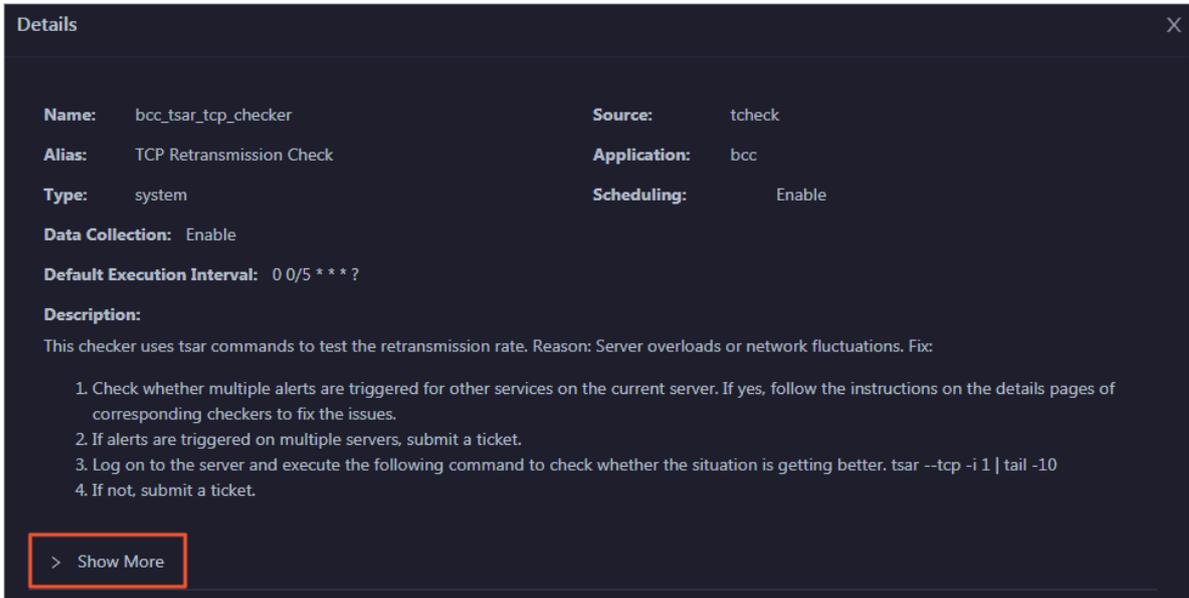
1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is

enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

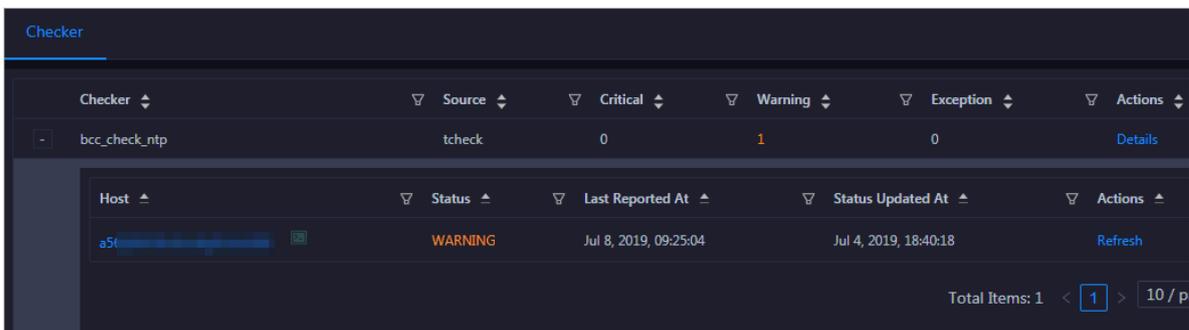


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

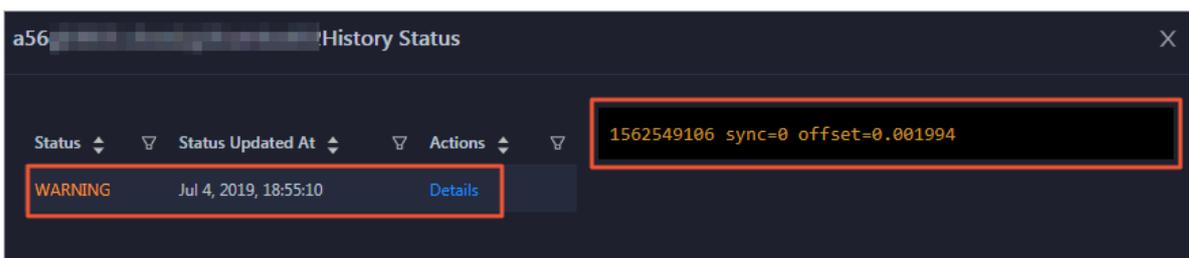
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

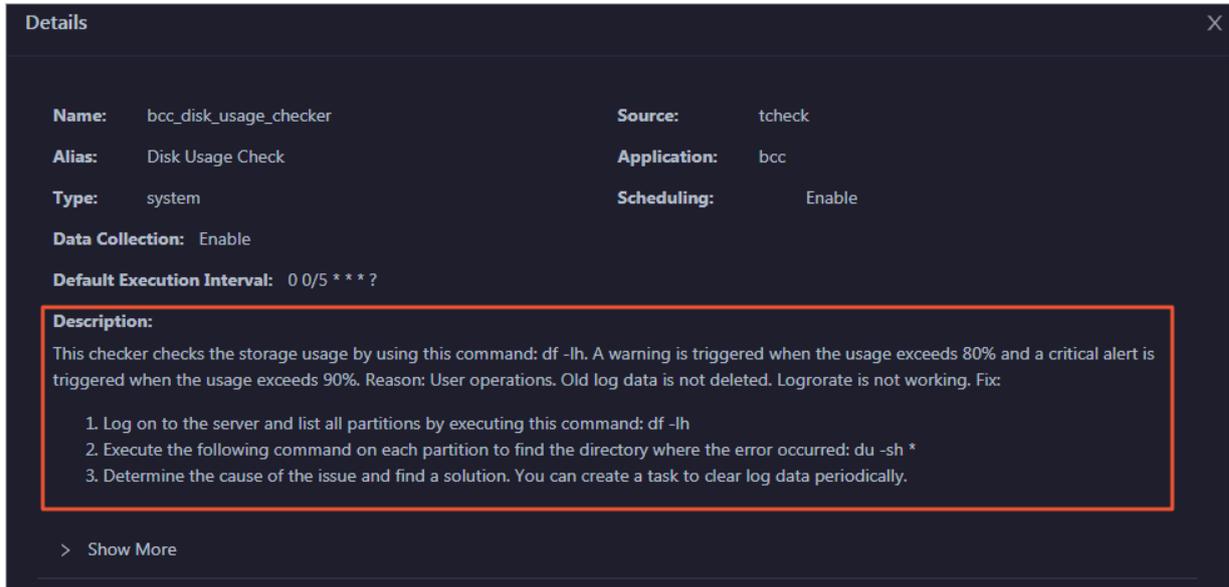


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

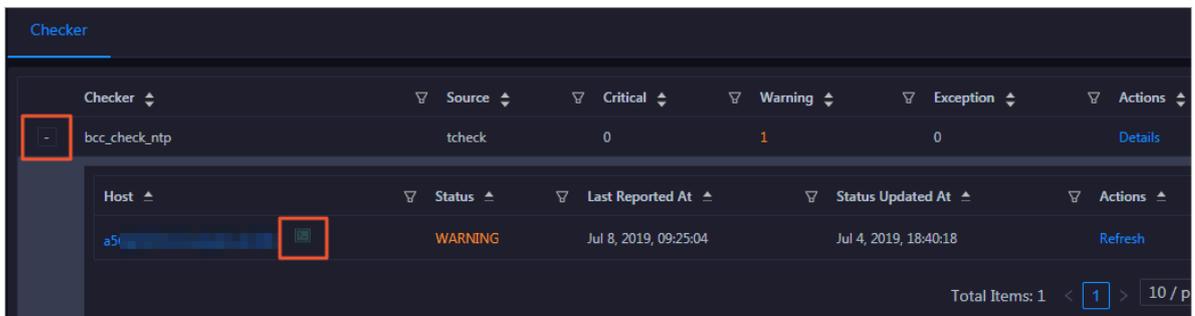
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



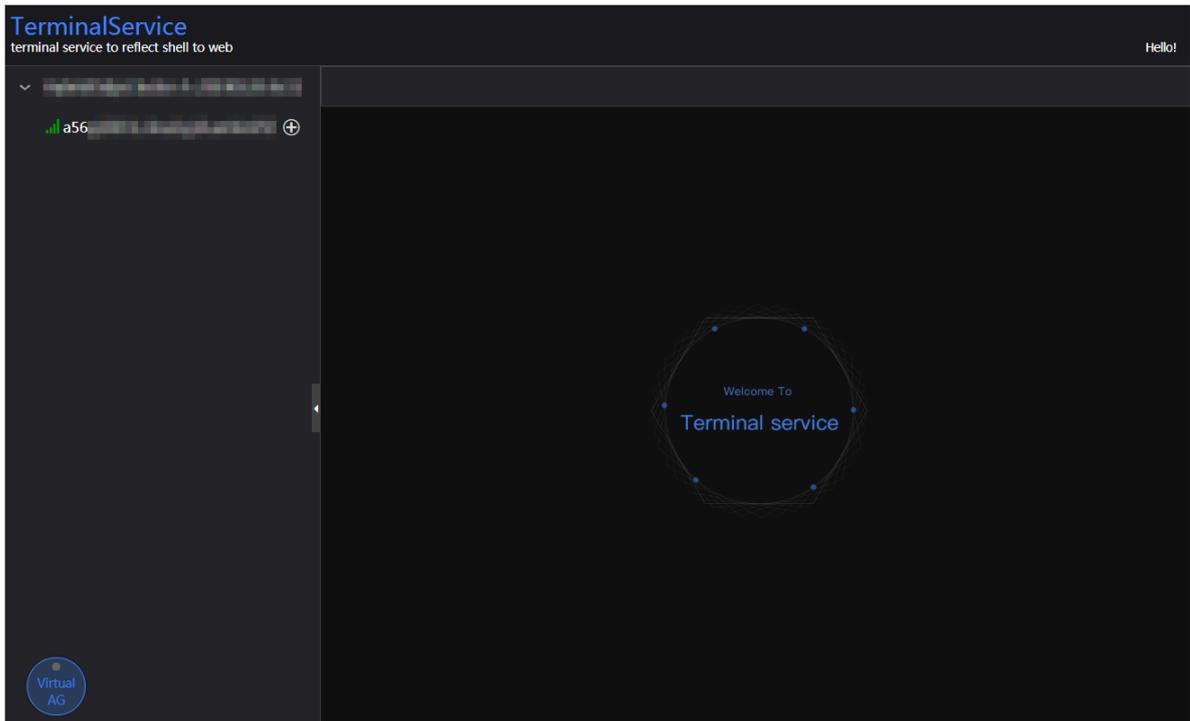
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

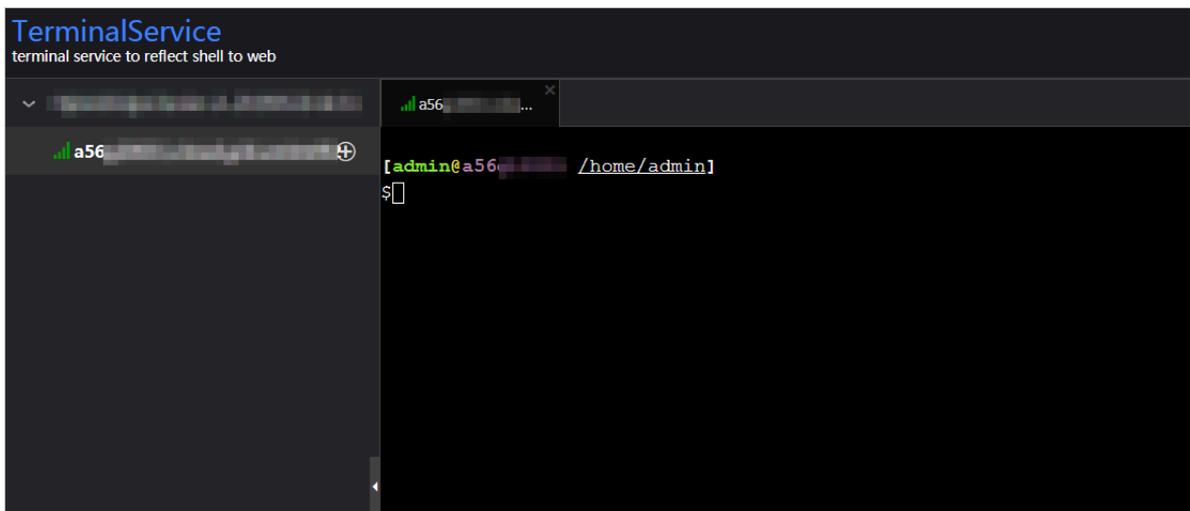
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

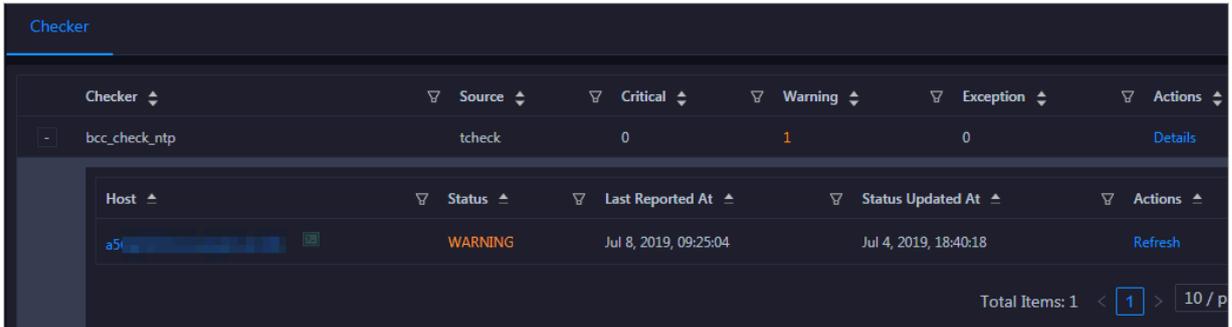


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

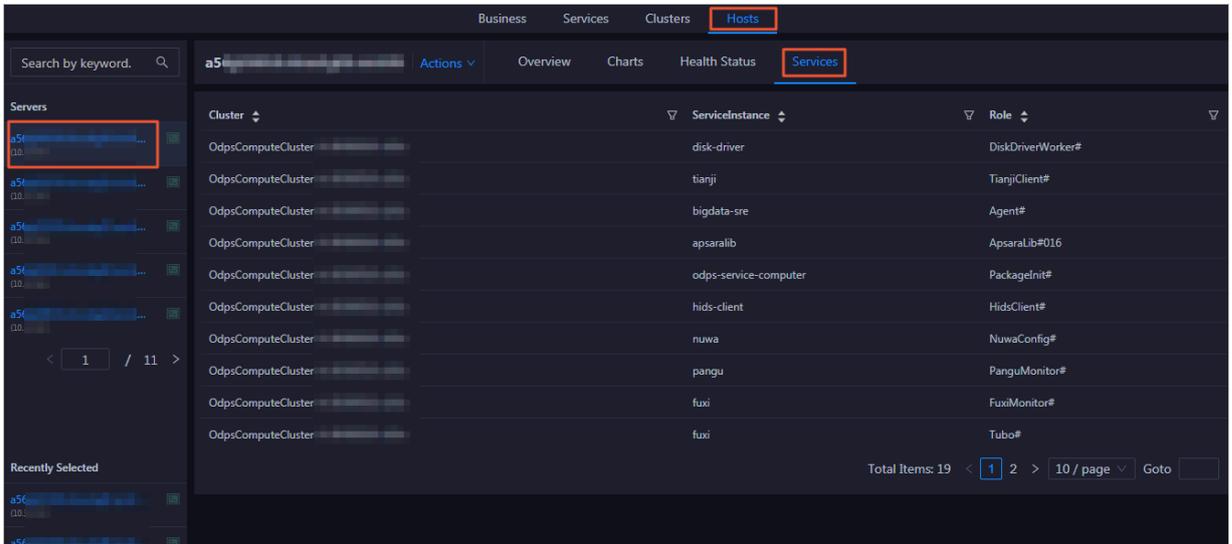
After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.5.5.5 Host services

On the Services page, you can view information about service instances and service instance roles of a host.

On the Hosts page, select a host in the left-side navigation pane, and then click the Services tab. The Services page for the host appears.



On the Services page, you can view the cluster, service instances, and service instance roles of the host.

3.1.6 StreamCompute

3.1.6.1 O&M overview

This topic describes the O&M features of Realtime Compute and how to access the Realtime Compute O&M page.

Modules

Realtime Compute O&M includes four modules including business O&M, service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Business O&M	Projects	Displays information about all projects in Realtime Compute.
	Jobs	Displays information about all jobs in Realtime Compute, and supports job diagnosis and analysis.
	Queues	Displays information about all queues in Realtime Compute.
Service O&M	Druid	Displays the number of Druid master nodes and that of Druid worker nodes in Realtime Compute .
	Yarn	Displays information about the YARN queue APIs in Realtime Compute. Realtime Compute allocates cluster resources to YARN queue APIs. You can bind projects with these APIs to obtain the corresponding cluster resources.
Cluster O&M	Overview	Displays the overall running and health check information about a cluster. On this page, you can view the health check result and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Health Status	Displays the check results for a cluster. The check results are divided into Critical, Warning, Exception, and OK.
	Hosts	
Host O&M	Overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Health Status	Displays the check results for a host. The check results are divided into Critical, Warning, Exception, and OK.

Entry

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click StreamCompute.
3. On the StreamCompute page that appears, click O&M in the upper-right corner.
The Business page appears.

The O&M page includes four modules, namely, Business, Services, Clusters, and Hosts.

3.1.6.2 Business O&M

3.1.6.2.1 Projects

Apsara Bigdata Manager (ABM) allows you to view information about the projects in Realtime Compute, including the name, engine, queue, used CUs, total CUs, CU usage percentage, and number of jobs.

On the Business page, click Projects in the left-side navigation pane. The Projects page for Realtime Compute appears.

The Projects page displays the information of projects in Realtime Compute, including the name, engine, queue, used CUs, total CUs, CU usage percentage, and number of jobs.

3.1.6.2.2 Jobs

Apsara Bigdata Manager (ABM) allows you to view the information about jobs in Realtime Compute, including the name, user, project, transactions per second (TPS) in the inbound direction, latency, requested CUs, status, and start time. You can diagnose and analyze jobs to troubleshoot issues.

Jobs

On the Business page, click Jobs in the left-side navigation pane. The Jobs page for Realtime Compute appears.

The Jobs page displays the information about jobs in Realtime Compute, including the name, user, project, TPS in the inbound direction, latency, requested CUs, status, and start time.

Job analysis

Job diagnosis has two steps, namely Failover and Blink Metric. In the Blink Metric step, the system checks the latency, garbage collection (GC) time, TPS, the number of times of GC, data skew, and back pressure nodes of a job.

- 1. On the Jobs page, click a job name. Alternatively, click the Job Analysis tab at the top. The Job Analysis page appears.**
- 2. Select the job to be diagnosed and analyzed from the Select Job drop-down list.**
- 3. In the Diagnosis section, click Start Diagnosis.**

After the diagnosis starts, the system automatically evaluates the time required for the diagnosis. Wait until the diagnosis is complete.

- 4. After the diagnosis is completed, click View Log to view the log details if the diagnosis result appears in red.**

The metrics for job diagnosis are described as follows:

- **Failover**
 - **Checks whether a failover is triggered for a job in a specified period and displays the information about the failover.**
- **Blink Metric**
 - **Job Latency: checks whether the latency of a subtask exceeds 10 minutes.**
 - **Job GC Time: checks whether the GC time of CMS exceeds 100 ms. This metric applies to all containers.**
 - **Job TPS: checks whether the TPS of a subtask is 0.**
 - **Number of GC Times: checks whether the number of the GC times exceeds 15 per minute. This metric applies to all containers.**
 - **Data Skew: checks whether the deviation of the input data size of each subtask in a task to the average input data size of all subtasks in the task exceeds 30%.**
 - **Back Pressure Nodes: checks whether each task has back pressure and finds the nodes that cause back pressure.**

3.1.6.2.3 Queues

Apsara Bigdata Manager (ABM) allows you to view the information about the queues in Realtime Compute, including the name, status, minimum resources

guaranteed, minimum resources guaranteed, maximum resources available, maximum resources available, and number of jobs.

On the Business page, click Queues in the left-side navigation pane. The Queues page for Realtime Compute appears on the right.

The Queues page displays the information about queues in Realtime Compute, including the name, status, minimum resources guaranteed, minimum resources guaranteed, maximum resources available, maximum resources available, and number of jobs.

3.1.6.3 Service O&M

3.1.6.3.1 Yarn

Apsara Bigdata Manager (ABM) allows you to view the overview and health status of the YARN service in Realtime Compute.

Overview

On the Services page, click Yarn in the left-side navigation pane. The Overview page for the YARN service appears.

The Overview page displays the health check result, health check history, application status, container status, node status, logical CPU usage, and logical memory usage for the YARN service.

Click View Details in the Health Check or Health Check History section. The Health Status page for the YARN service appears. On this page, you can view more details about the health check.

Heath status

On the Services page, click Yarn in the left-side navigation pane. Click the Health Status tab at the top of the Services page. The Health Status page for the YARN service appears.

On the Health Status page, you can view all checkers of the YARN service and the check results for all hosts. The check results are divided into Critical, Warning, Exception, and OK. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

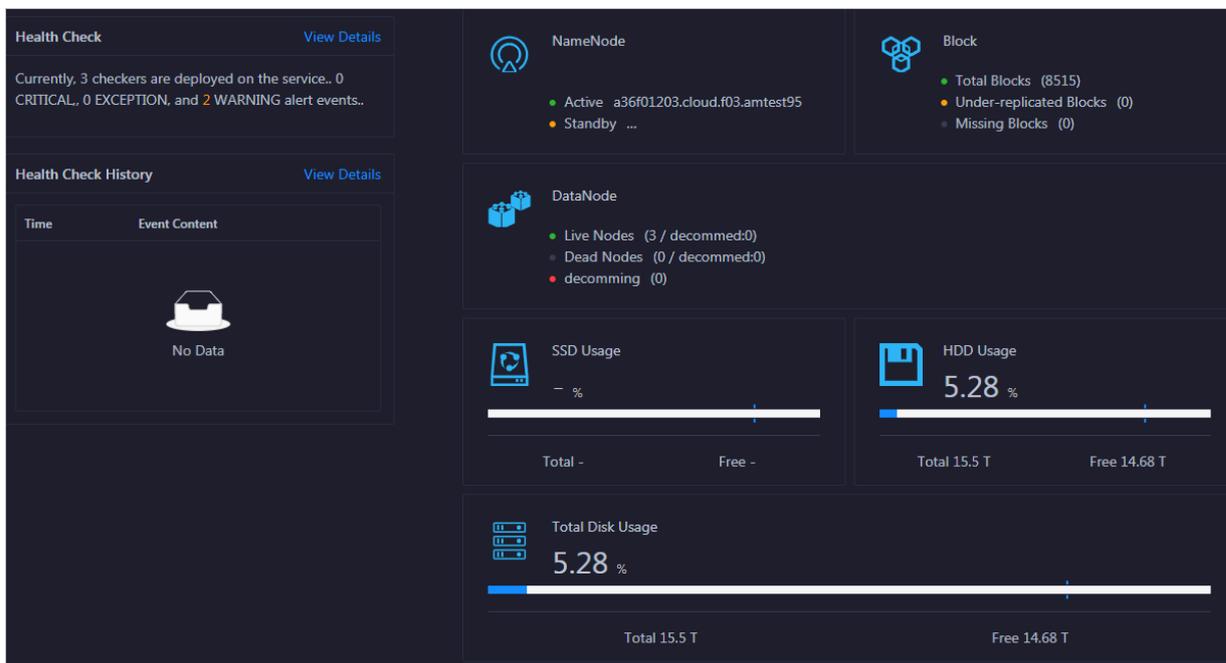
The operations you can perform on the Health Status page for the YARN service are the same as those on the Health Status page for Realtime Compute clusters. For more information, see [Cluster health](#).

3.1.6.3.2 HDFS

Apsara Bigdata Manager (ABM) allows you to view the overview and health status of the Hadoop Distributed File System (HDFS) service in Realtime Compute.

Overview

On the Services page, click HDFS in the left-side navigation pane. The Overview page for the HDFS service appears.



The Overview page displays the health check result, health check history, the information of NameNode, blocks, and DataNode, solid-state disk (SSD) usage, hard disk drive (HDD) usage, and total disk usage.

Click [View Details](#) in the Health Check section and [View Details](#) in the Health Check History section to go to the Health Status page for the HDFS service. On this page, you can view more details about the health check.

Health Status

On the Services page, click HDFS in the left-side navigation pane. Click the Health Status tab. The Health Status page for the HDFS service appears.

Checker	Source	Critical	Warning	Exception	Actions										
- streamcompute_HDFS_FilesAndBlockTotal_checker	tcheck	0	1	0	Details										
<table border="1"> <thead> <tr> <th>Host</th> <th>Status</th> <th>Last Reported At</th> <th>Status Updated At</th> <th>Actions</th> </tr> </thead> <tbody> <tr> <td></td> <td>WARNING</td> <td>Dec 9, 2019, 16:30:02</td> <td>Nov 22, 2019, 16:45:03</td> <td>Refresh</td> </tr> </tbody> </table>						Host	Status	Last Reported At	Status Updated At	Actions		WARNING	Dec 9, 2019, 16:30:02	Nov 22, 2019, 16:45:03	Refresh
Host	Status	Last Reported At	Status Updated At	Actions											
	WARNING	Dec 9, 2019, 16:30:02	Nov 22, 2019, 16:45:03	Refresh											
+ streamcompute_HDFS_CapacityUsed_checker	tcheck	0	1	0	Details										
+ streamcompute_HDFS_checker	tcheck	0	0	0	Details										

On the Health Status page, you can view all checkers of the HDFS service and the check results for all hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

The operations you can perform on the Health Status page for the HDFS service are the same as those on the Health Status page for Realtime Compute clusters. For more information, see [Cluster health](#).

3.1.6.4 Cluster O&M

3.1.6.4.1 Cluster overview

The cluster overview page displays the overall running and health check information about a cluster. On this page, you can view the health check result and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Entry

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.

The Overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the cluster. On this page, you can also view the health check result and health check history of the cluster. To view information about a cluster, select a region in the left-side navigation pane, and then select the cluster in the region.

Hosts

This section displays all host statuses and the number of hosts in each status. The host statuses include good and bad.

Services

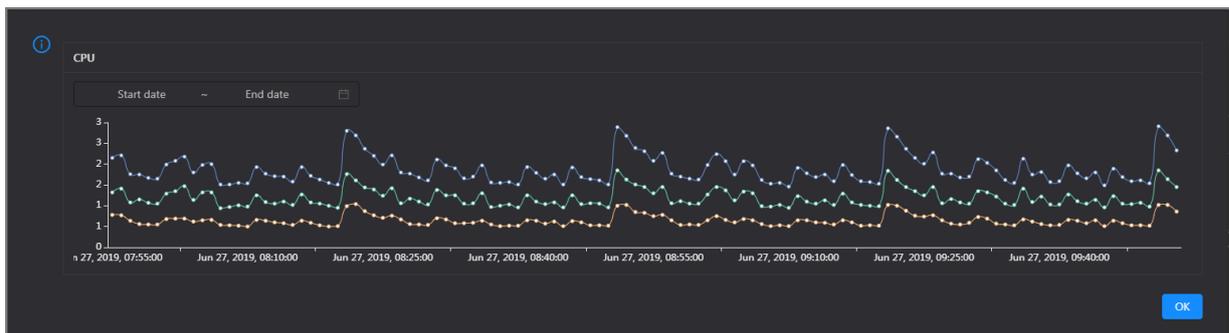
This section displays all services deployed in the cluster and the respective number of available and unavailable services.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

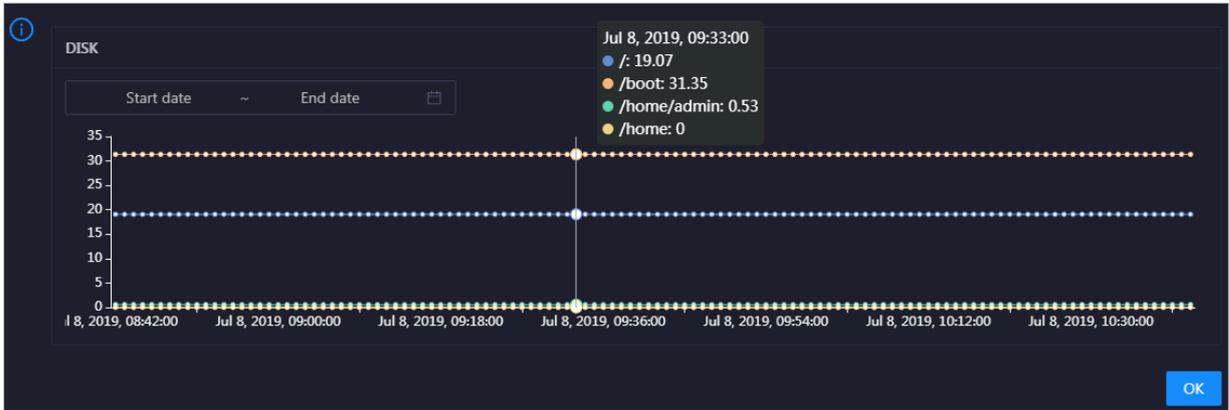
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

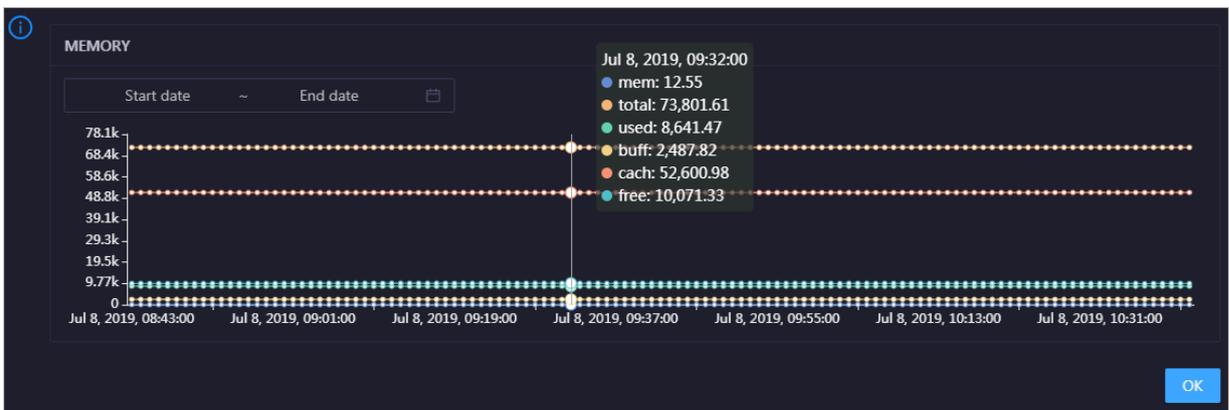


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

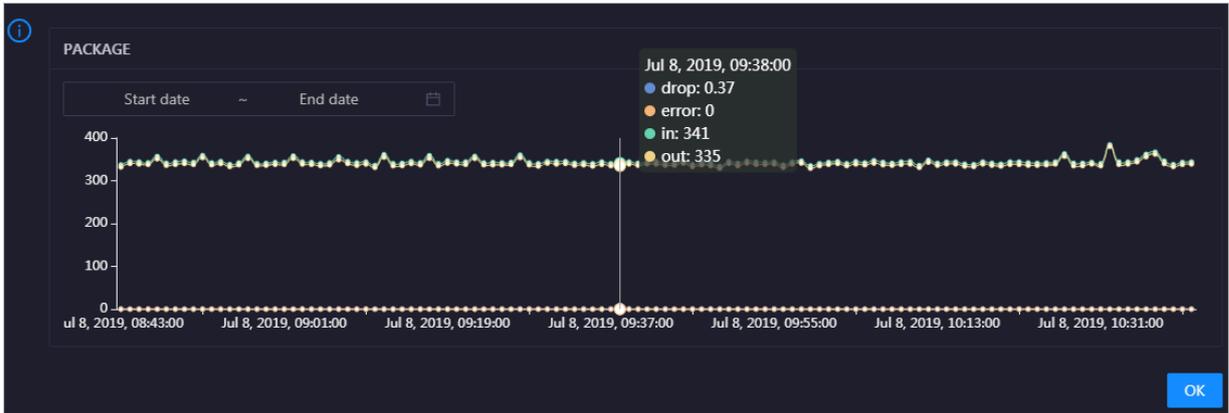


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

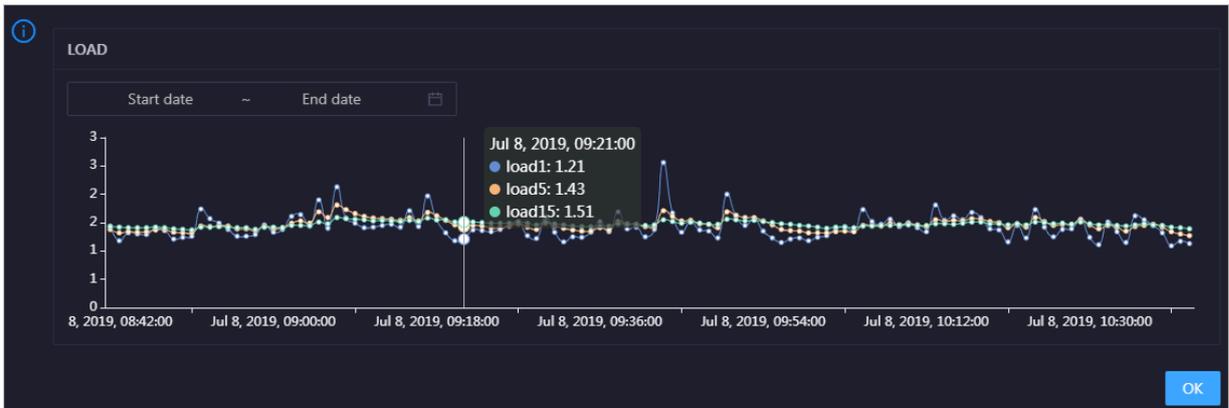


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

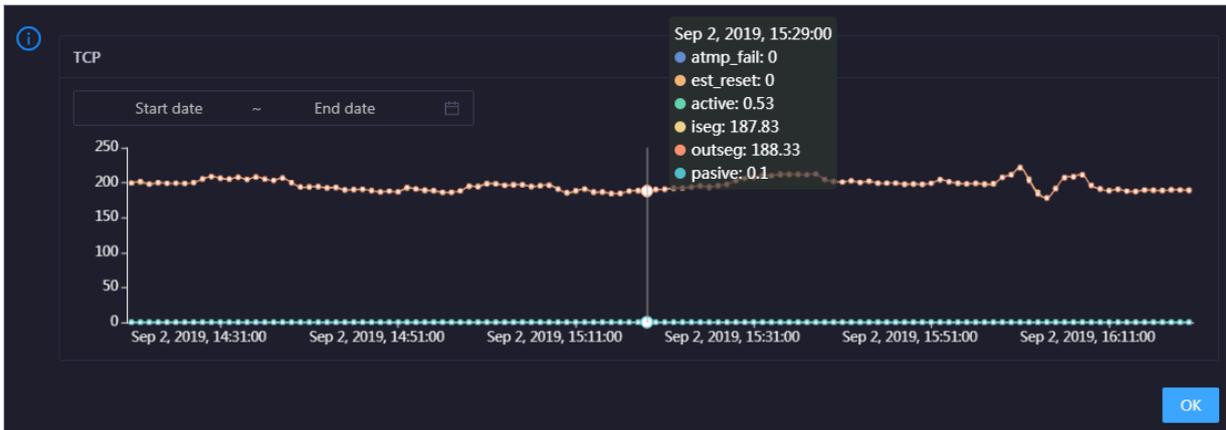


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

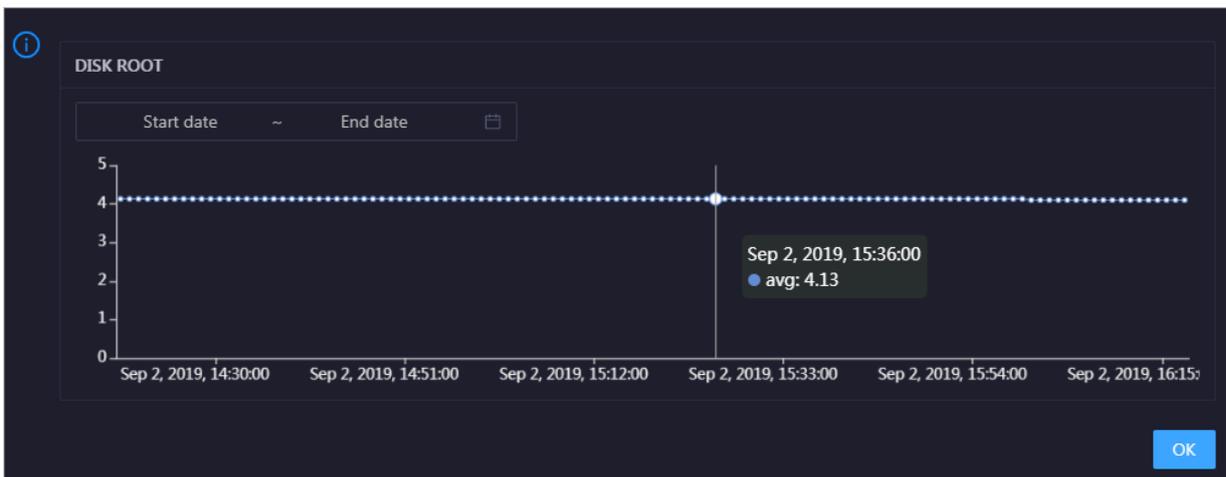


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

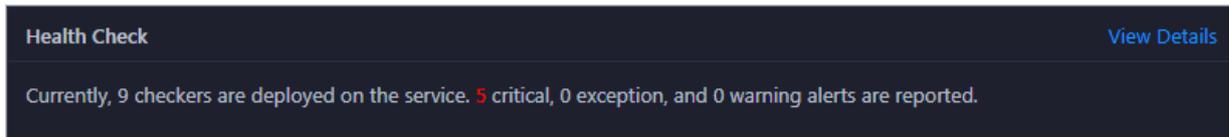
Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

Health Check

This section displays the number of checkers deployed for the cluster and the respective number of Critical, Warning, and Exception alerts.



Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

Health Check History

This section displays a record of the health checks performed on the cluster.

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

You can click the event content of a check to view the exception items.

3.1.6.4.2 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

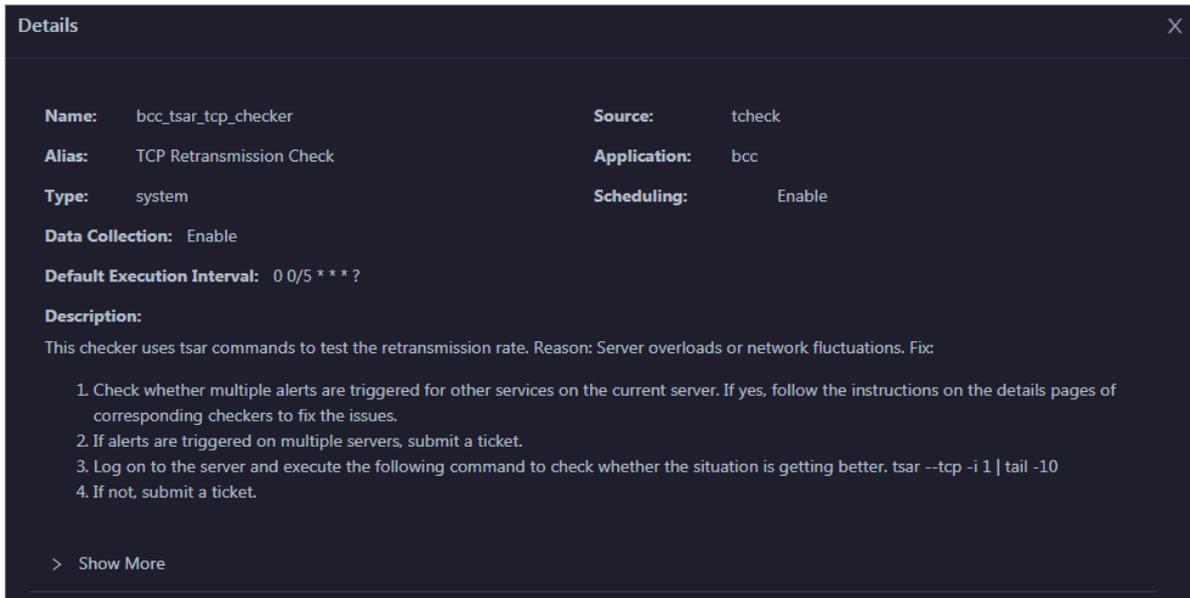
Entry

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

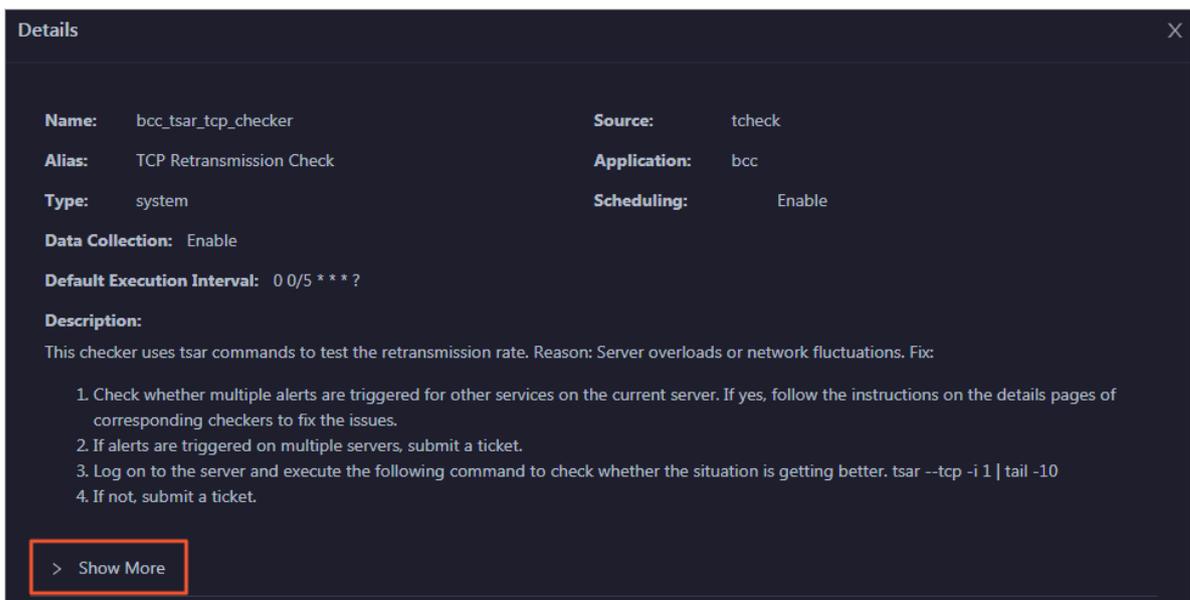
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

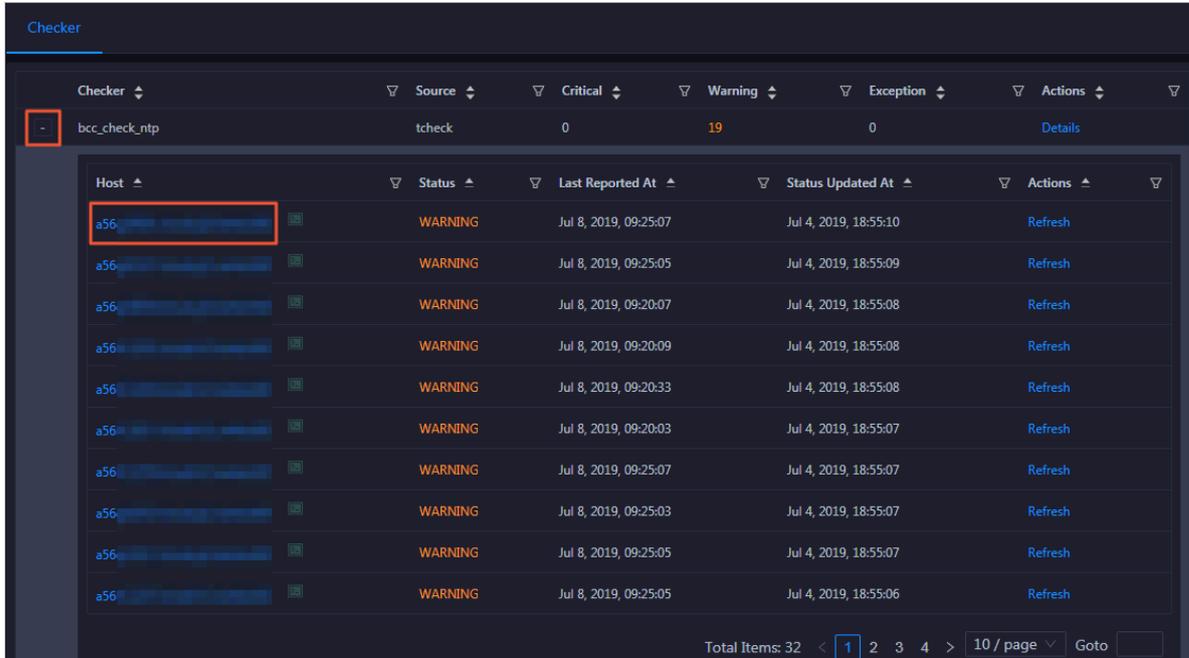


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

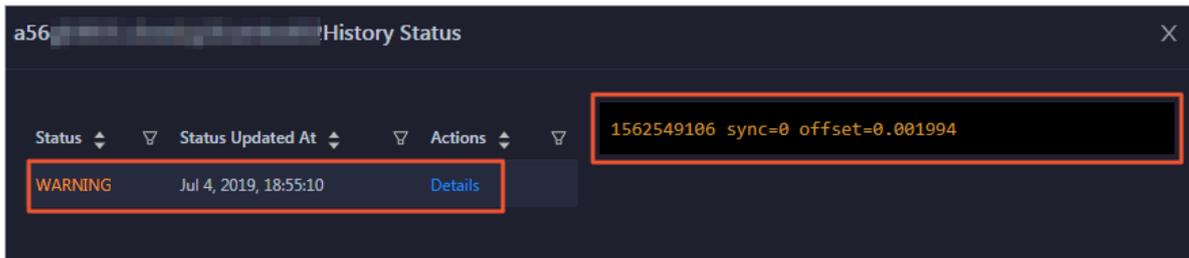
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

- 1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.**



- 2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**



Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

Details ✕

Name: bcc_disk_usage_checker	Source: tcheck
Alias: Disk Usage Check	Application: bcc
Type: system	Scheduling: Enable
Data Collection: Enable	
Default Execution Interval: 0 0/5 * * * ?	

Description:
 This checker checks the storage usage by using this command: `df -lh`. A warning is triggered when the usage exceeds 80% and a critical alert is triggered when the usage exceeds 90%. Reason: User operations. Old log data is not deleted. Logrotate is not working. Fix:

1. Log on to the server and list all partitions by executing this command: `df -lh`
2. Execute the following command on each partition to find the directory where the error occurred: `du -sh *`
3. Determine the cause of the issue and find a solution. You can create a task to clear log data periodically.

> Show More

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

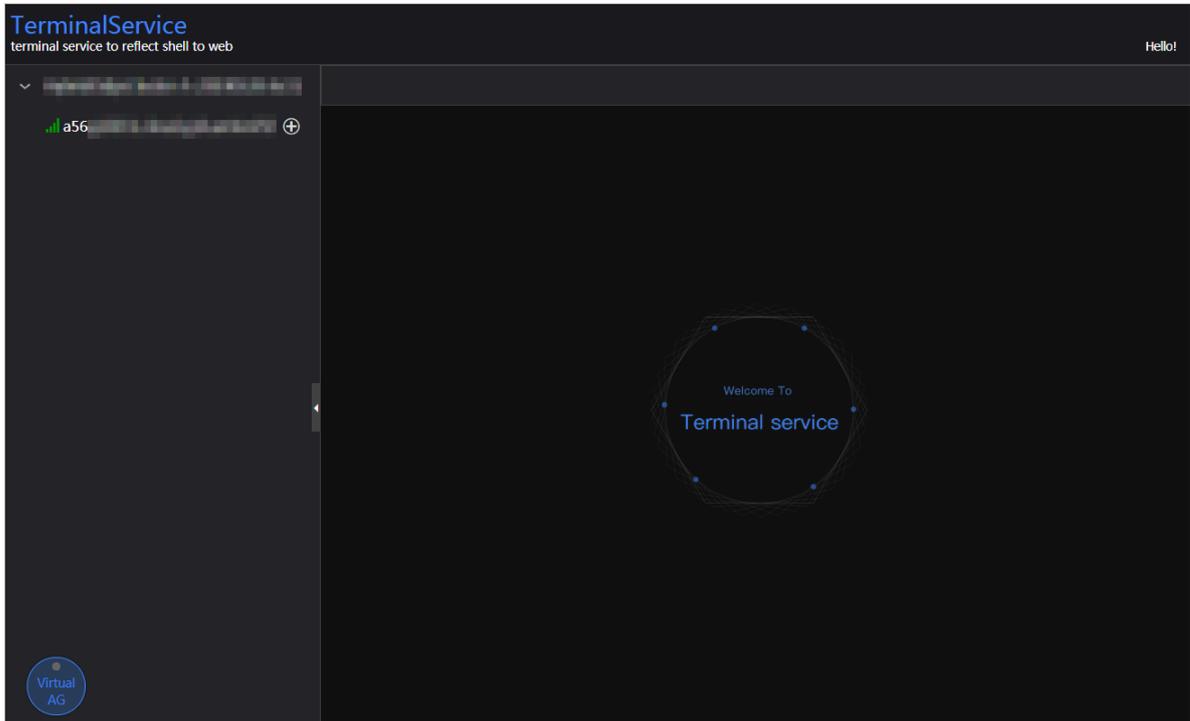
- 1. On the Health Status page, click + to expand a checker with alerts.**

Checker

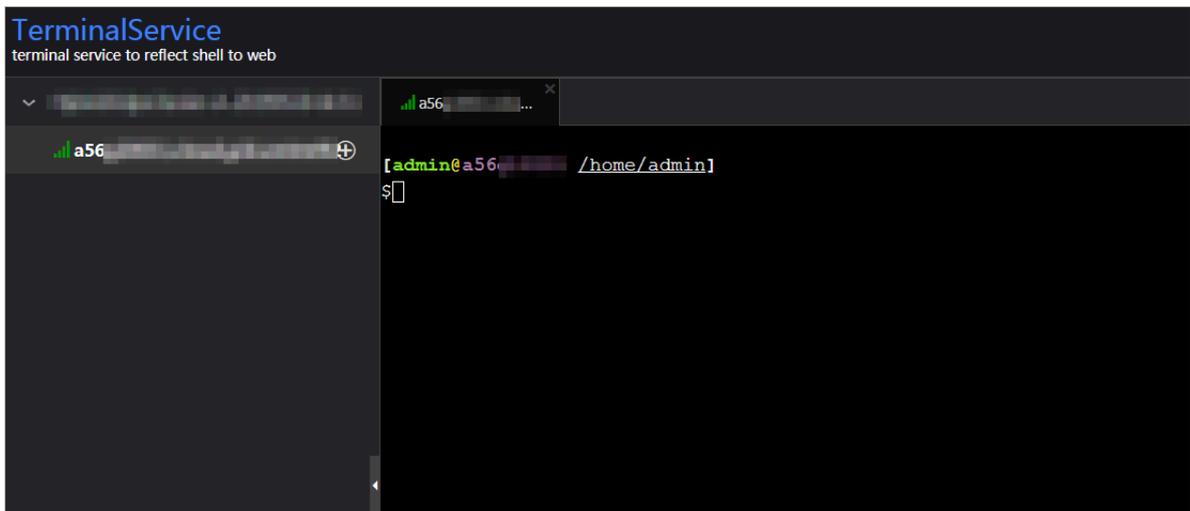
Checker	Source	Critical	Warning	Exception	Actions
- bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56 +	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56 +	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56 +	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.



3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

The screenshot shows a 'Checker' interface with a table of hosts. The table has columns for Host, Status, Last Reported At, Status Updated At, and Actions. The first row shows a host with a 'WARNING' status and a 'Refresh' button highlighted in a red box.

Host	Status	Last Reported At	Status Updated At	Actions
a56...	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56...	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56...	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56...	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56...	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56...	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56...	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

3.1.6.4.3 Hosts

The Hosts page displays information about hosts, including the hostname, IP address, role, type, CPU usage, total memory size, available memory size, load, root disk usage, packet loss rate, and packet error rate.

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Hosts tab. The Hosts page for the cluster appears.

To view more information about a host, click the name of the host. The Overview tab of the Hosts page appears. For more information, see [Host overview](#).

3.1.6.4.4 Cluster scale-out

Apsara Bigdata Manager (ABM) allows you to scale out a Realtime Compute cluster by adding physical hosts. Cluster scale-out refers to the process of adding physical hosts in the default cluster of Apsara Infrastructure Management Framework to a Realtime Compute cluster. Currently, scale-out is only available for worker nodes in a Realtime Compute cluster.

Prerequisites

- Your ABM account must have the required permissions to perform O&M operations on Realtime Compute.
- The default cluster of Apsara Infrastructure Management Framework has hosts whose product type is blink.

Context

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to

the default cluster of Apsara Infrastructure Management Framework. The default cluster of Apsara Infrastructure Management Framework can be considered as an idle resource pool that provides resources for scaling out clusters for your business. To scale out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

Step 1: Obtain the name of the host to be added to a Realtime Compute cluster

Before the scale-out, obtain the name of the host in the default cluster of Apsara Infrastructure Management Framework.

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click TIANJI to log on to the Apsara Infrastructure Management Framework console.
3. In the top navigation bar of the page that appears, choose Operations > Machine Operations.
4. On the Machine Operations page that appears, search for a host whose product type is blink in the default cluster. Copy the name of the host.

Step 2: Add the host to a Realtime Compute cluster

You can add multiple hosts to a Realtime Compute cluster at a time to scale out the cluster. To achieve this, you need to first specify an existing host as the template host. When you scale out the Realtime Compute cluster, the hosts copy configurations from the template host so that the hosts can be added to the cluster at a time.

1. On the O&M page of the ABM console, click the Clusters tab. On the page that appears, select a cluster in the left-side navigation pane. Click the Hosts tab, and then select a host whose role is Worker as the template host.

2. **Choose Actions > Scale out Cluster.** In the Scale out Cluster dialog box that appears, set relevant parameters.

The parameters are described as follows:

- **Refer Hostname:** the name of the template host. By default, the name of the selected host is used.
- **Hostname:** the name of the host to be added to the Realtime Compute cluster. The drop-down list displays all available hosts in the default cluster for scale-out. You can select one or more hosts from the drop-down list.

3. **Click Run.** A message appears, indicating that the action has been submitted.
4. **View the scale-out status.**

Click **Actions** in the upper-left corner, and then click **Execution History** next to **Scale out Cluster** to view the scale-out history.

It may take some time for the cluster to be scaled out. In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

Step 3: View the scale-out progress

If the status is **RUNNING**, click **Details** to view the steps and progress of the scale-out.

(Optional) Step 4: Locate the cause of a scale-out failure

If the status is **FAILED**, click **Details** to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

3.1.6.4.5 Cluster scale-in

Apsara Bigdata Manager (ABM) allows you to remove physical hosts to scale in a Realtime Compute cluster. Cluster scale-in refers to the process of removing physical hosts from a Realtime Compute cluster to the default cluster of Apsara Infrastructure Management Framework. Currently, scale-in is only available for the worker nodes in a Realtime Compute cluster.

Prerequisites

- **Your ABM account must have the required permissions to perform O&M operations on Realtime Compute.**

- The current cluster has more than three worker nodes. A Realtime Compute cluster creates three replicas for data by default. At least three worker nodes are required. Make sure that the cluster has at least three worker nodes after scale-in.
- Before you scale in a cluster, check whether the resources of the cluster, including the disk, CPU, and memory, are still sufficient if the cluster is scaled in. For more information about how to check CPU usage and memory usage, see [Yarn](#). You can run the `df` command to check disk usage.

**Notice:**

Scale-in triggers a job failover on hosts. If the cluster resources are insufficient after scale-in, the failover fails. This leads to negative effects on your business.

Context

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to the default cluster of Apsara Infrastructure Management Framework. The default cluster of Apsara Infrastructure Management Framework can be considered as an idle resource pool that provides resources for scaling out clusters for your business. ABM allows you to scale in or out a cluster for your business. To scale out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework. You can remove multiple hosts from a Realtime Compute cluster at a time to scale in the cluster.

Procedure

1. On the O&M page of the ABM console, click the Clusters tab. On the page that appears, select a cluster in the left-side navigation pane. Click the Hosts tab, and then select one or more hosts whose role is Worker.
2. On the Clusters page, choose Actions > Scale in Cluster. The Scale in Cluster dialog box appears.
Hostname: the name of the host to be removed from the Realtime Compute cluster. By default, the name of the selected host is used.
3. Click Run. A message appears, indicating that the action has been submitted.

4. View the scale-in status.

Click **Actions** in the upper-left corner, and then click **Execution History** next to **Scale in Cluster** to view the scale-in history.

It may take some time for the cluster to be scaled in. In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

5. View the scale-in progress.

If the status is **RUNNING**, click **Details** to view the steps and progress of the scale-in.

6. Locate the cause of a scale-in failure.

If the status is **FAILED**, click **Details** to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

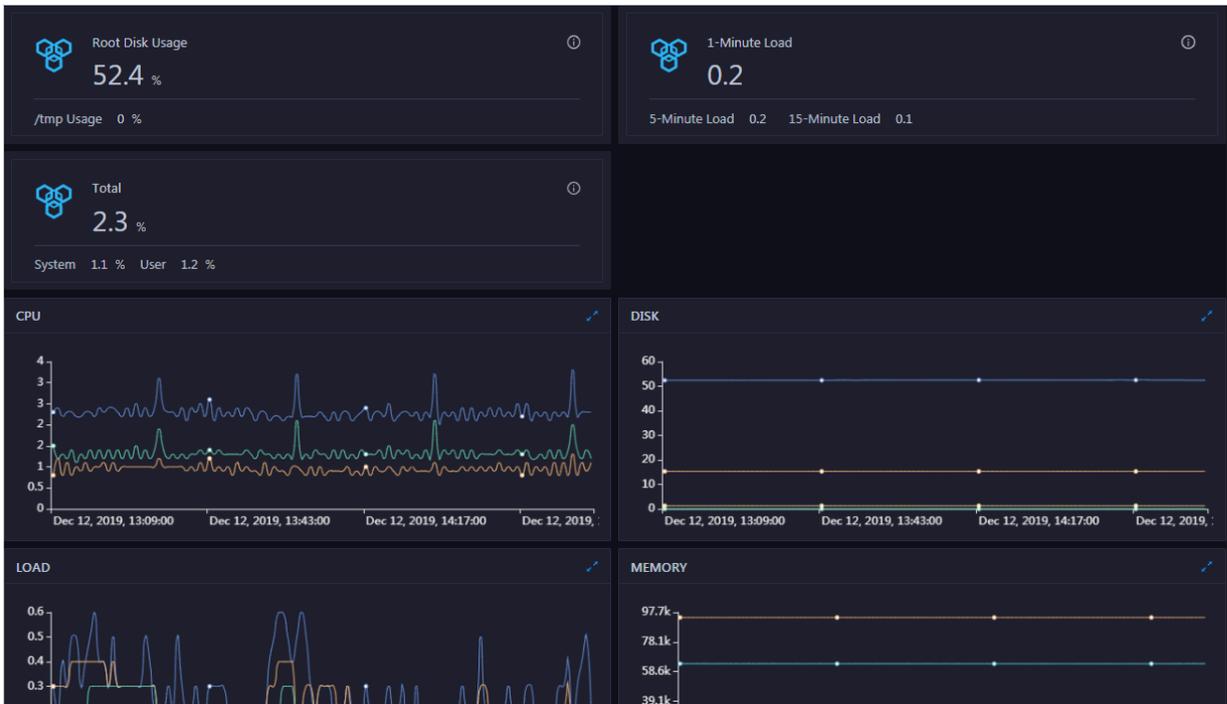
3.1.6.5 Host O&M

3.1.6.5.1 Host overview

The host overview page displays the overall running information about a host in a Realtime Compute cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Entry

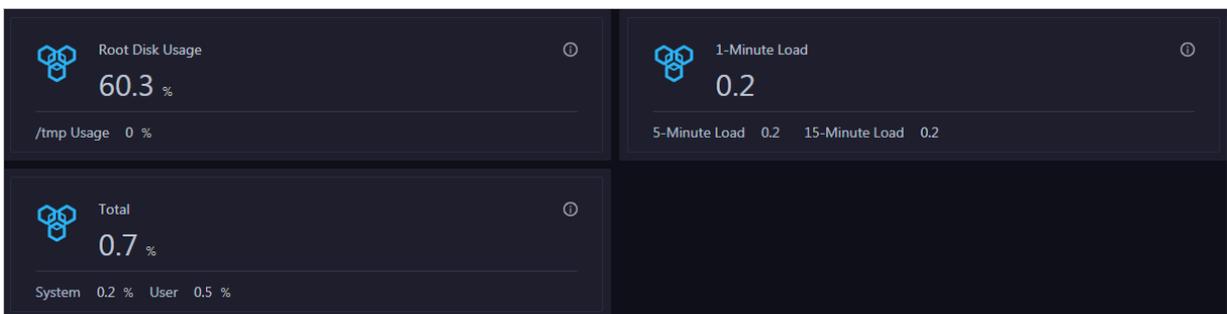
On the **Hosts** page, select a host in the left-side navigation pane. The **Overview** page for the host appears.



On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Root Disk Usage, Total, and 1-Minute Load

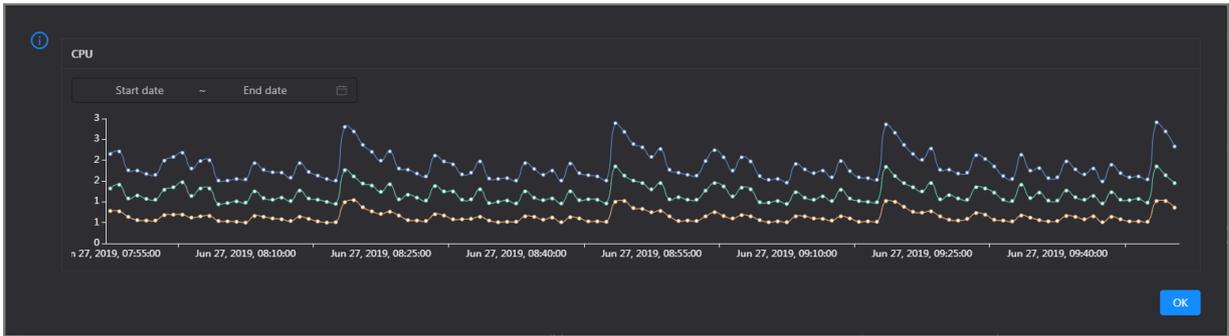
These sections display the root disk usage, total usage, and 1-minute load for the selected host. The Root Disk Usage section provides the usage of the /tmp directory. The Total section provides the system usage and user usage. The 1-Minute Load section provides the 1-minute, 5-minute, and 15-minute load averages.



CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

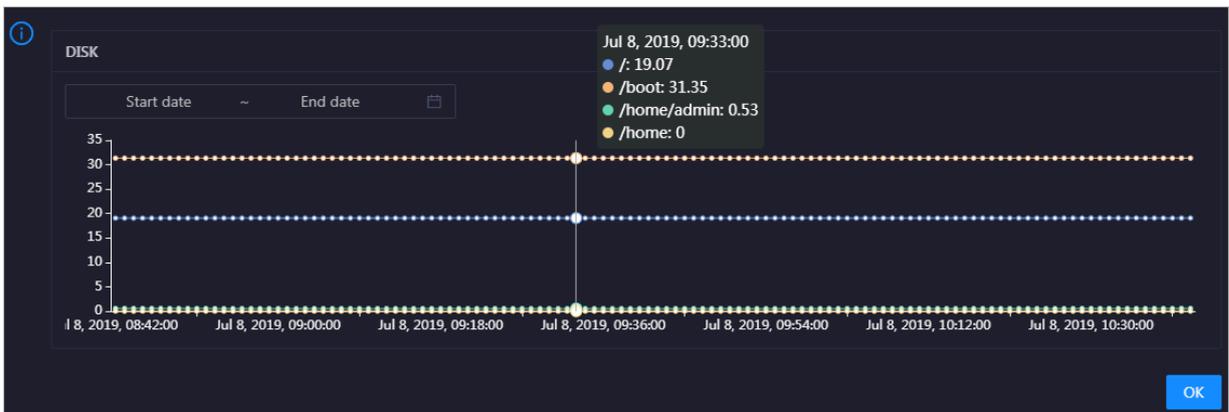


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

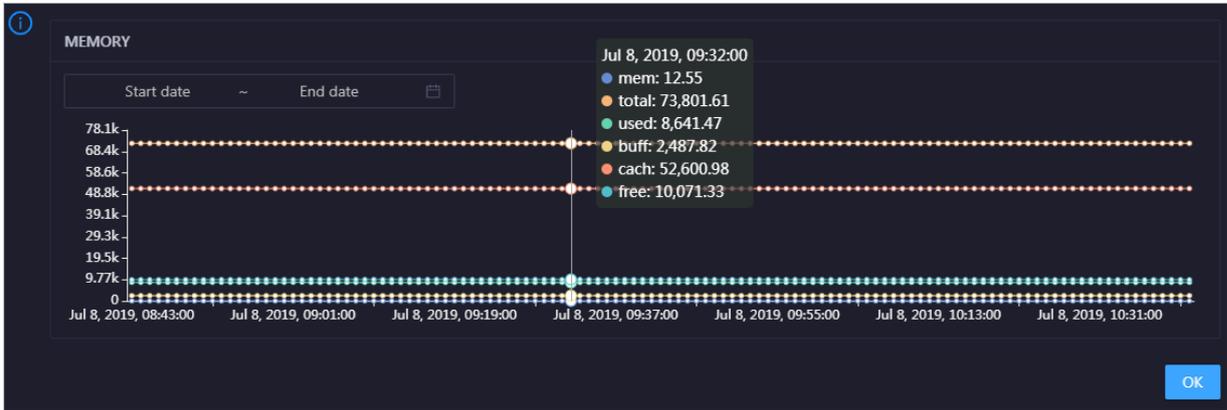


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (`mem`), total memory size (`total`), used memory size (`used`), size of memory used by kernel buffers (`buff`), size of memory used by the page cache (`cach`), and available memory size (`free`) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

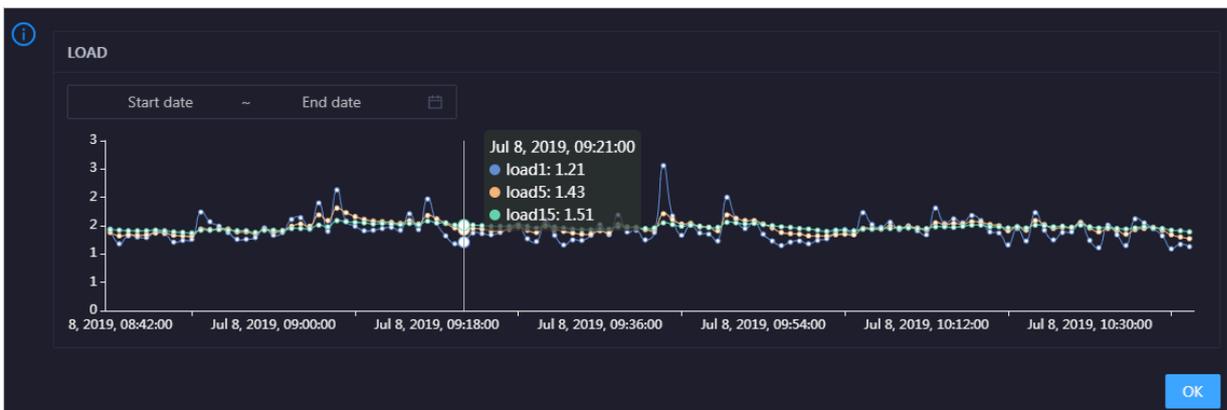


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

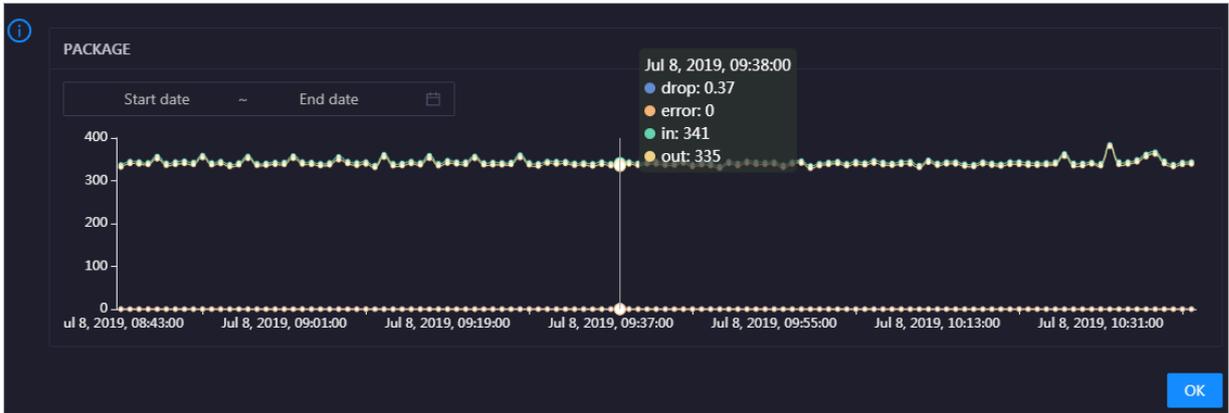


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.

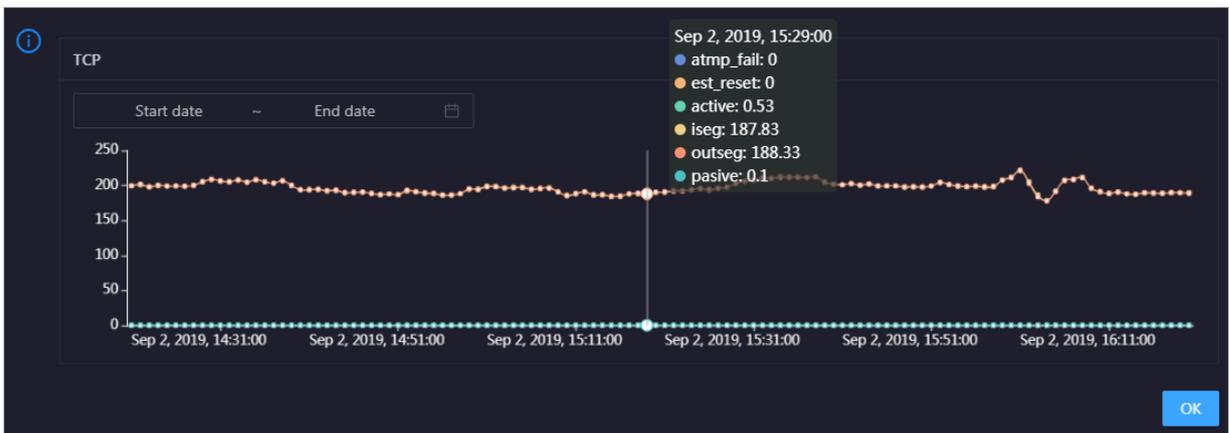


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click  in the upper-right corner of the chart to zoom in the chart.

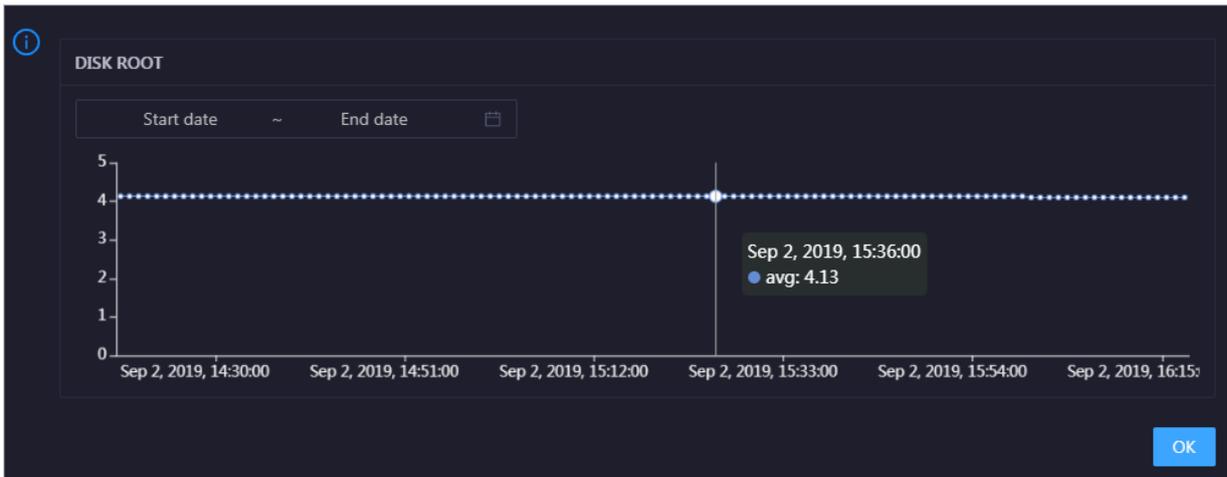


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the host over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

Health Check [View Details](#)

Currently, 10 checkers are deployed on the service. 0 critical, 0 exception, and 1 warning alerts are reported.

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

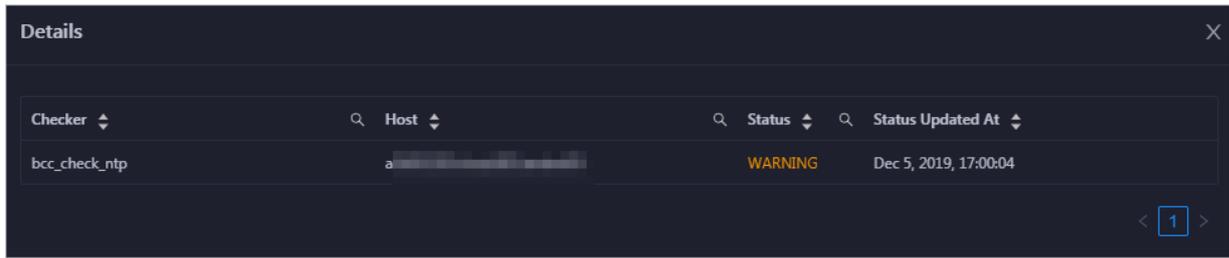
Health Check History

This section displays a record of the health checks performed on the host.

Health Check History		View Details
Time	Event Content	
Recently	1 alerts are reported by checkers.	< 1 >

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

You can click the event content of a check to view the exception items.



3.1.6.5.2 Host health

On the host health status page, you can view the checkers of all hosts, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

On the Hosts page, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.

The screenshot shows a table titled 'Checker' with the following columns and data:

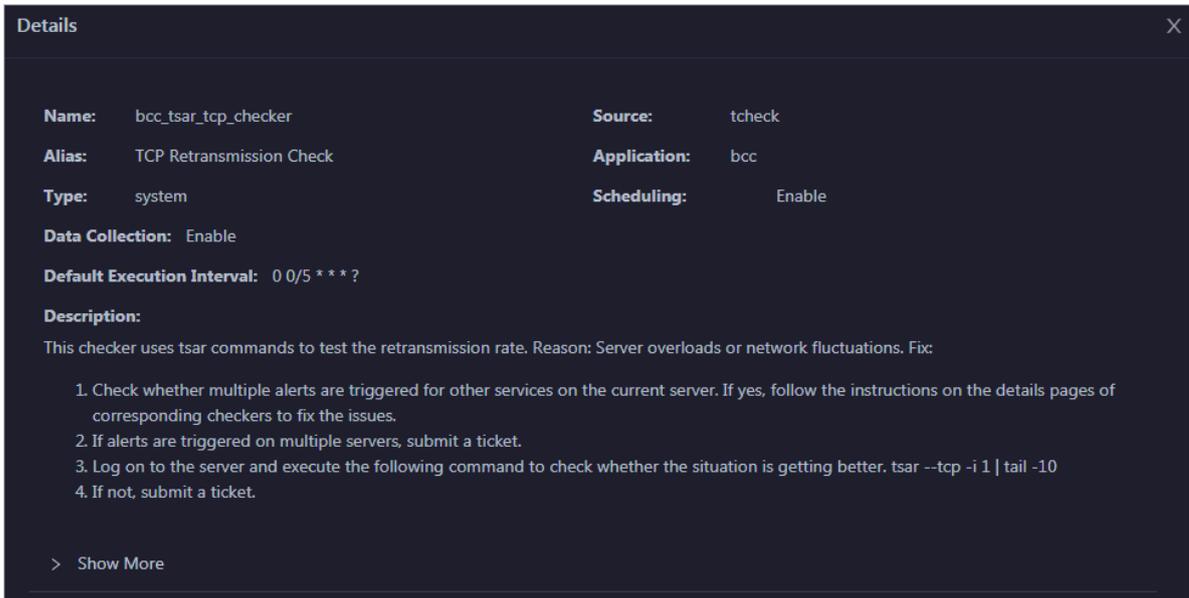
Checker	Source	Critical	Warning	Exception	Actions
+ bcc_check_ntp	tcheck	0	0	0	Details
+ bcc_disk_usage_checker	tcheck	0	0	0	Details

Navigation arrows and a page indicator '1' are visible at the bottom right.

On the Health Status page, you can view all checkers of the host and the check results for the hosts in the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

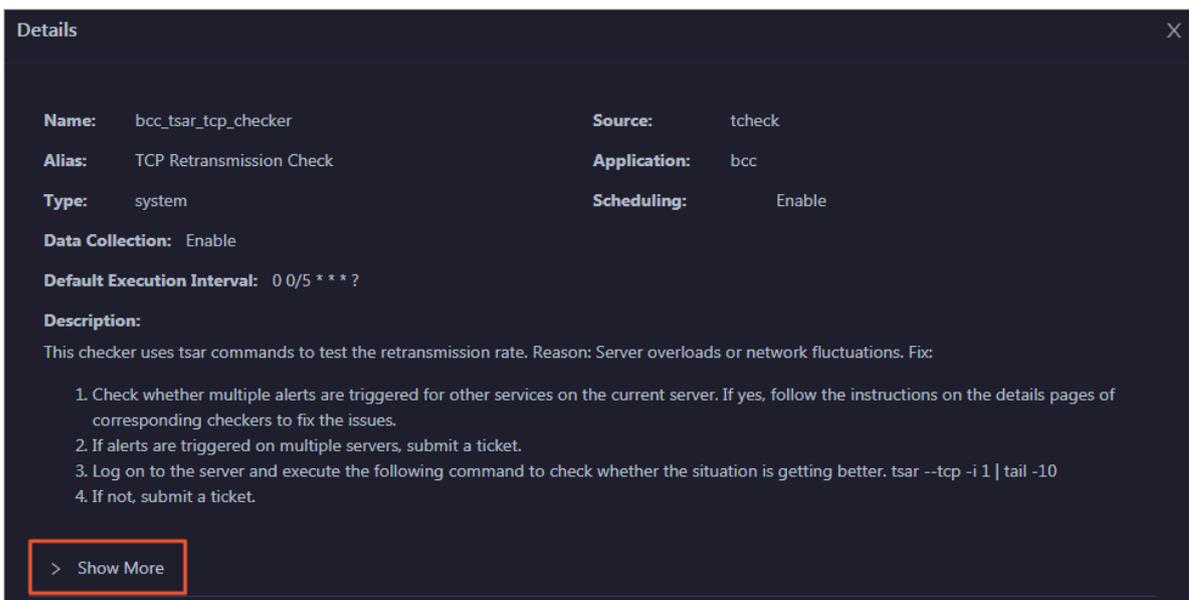
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

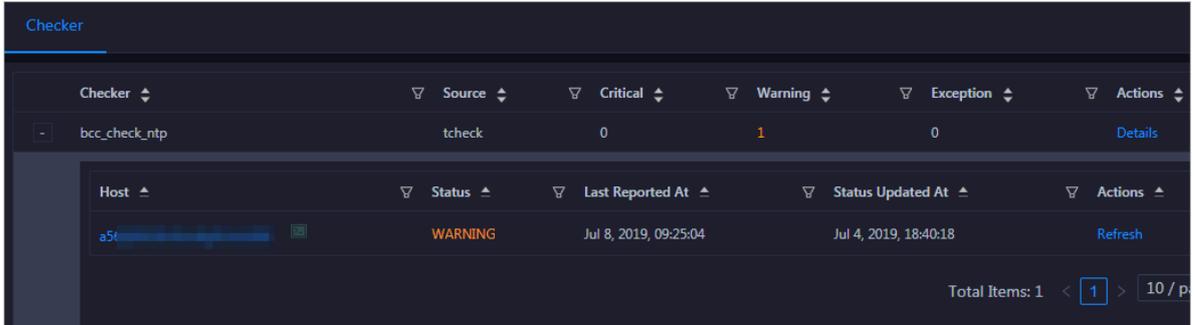


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

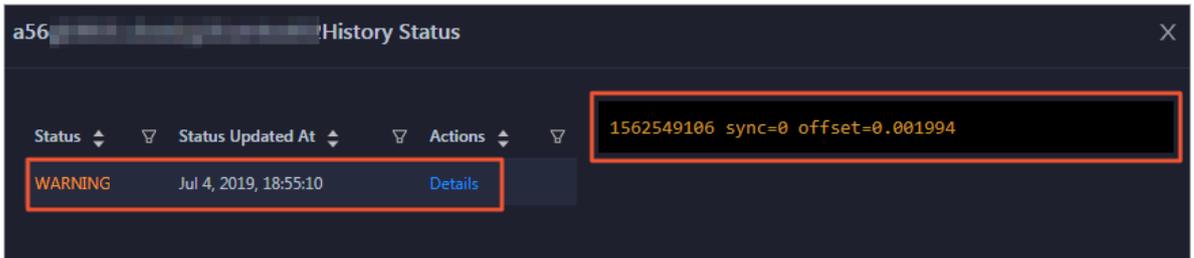
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

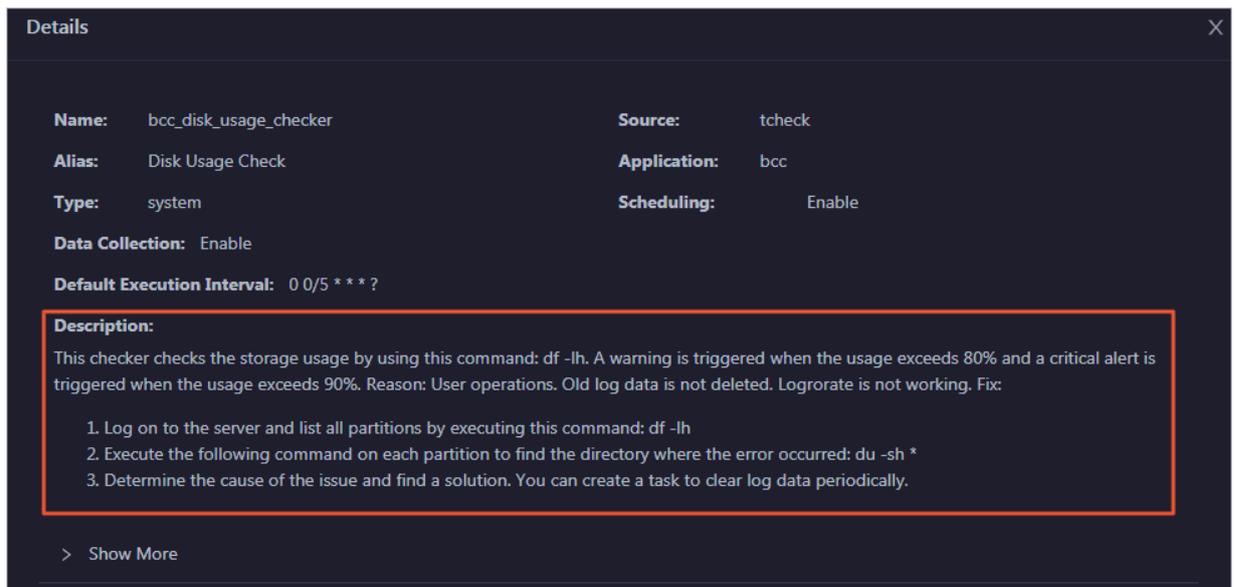


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

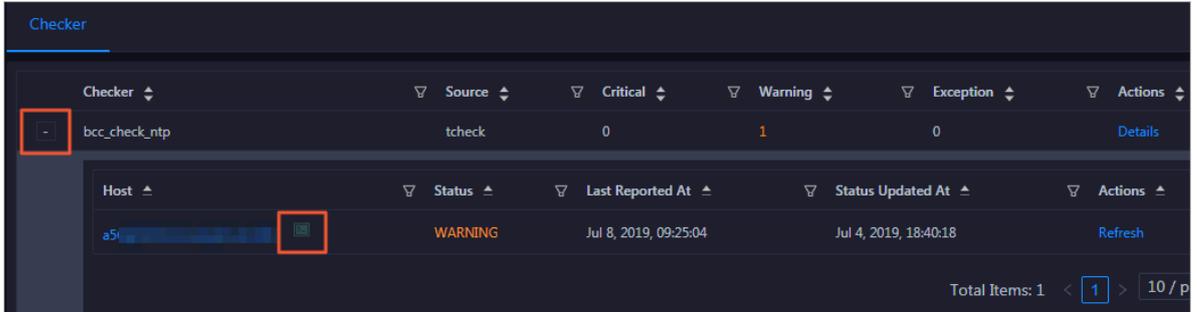
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



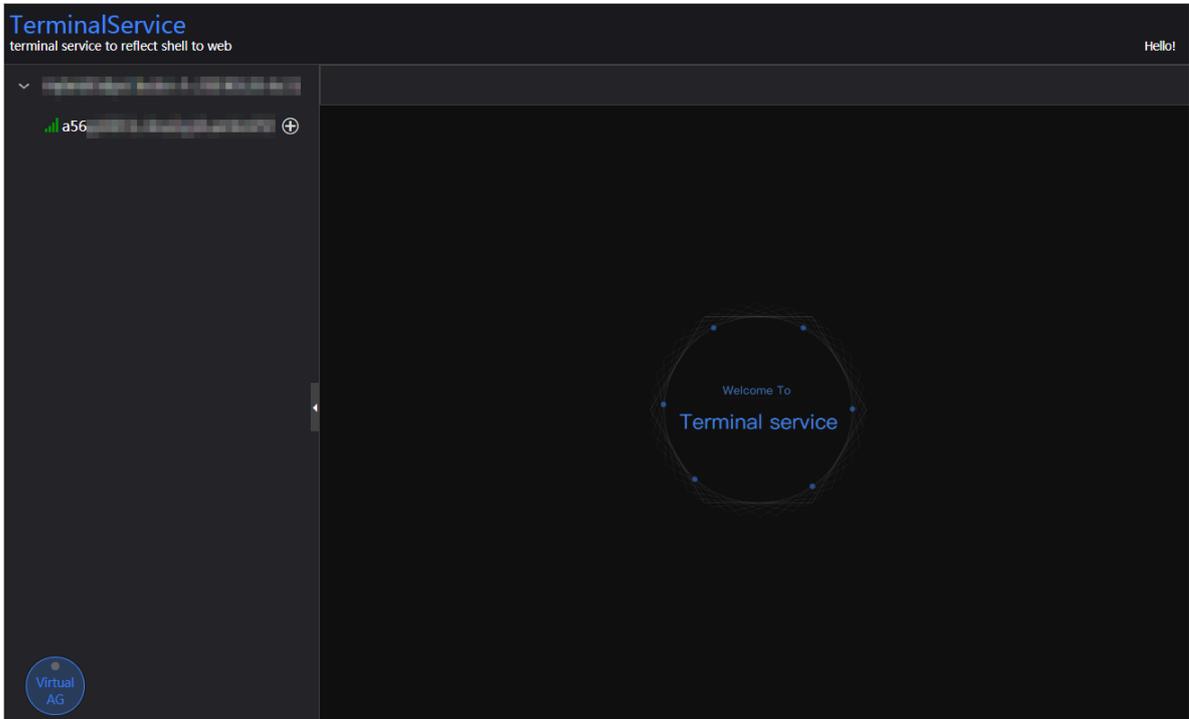
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

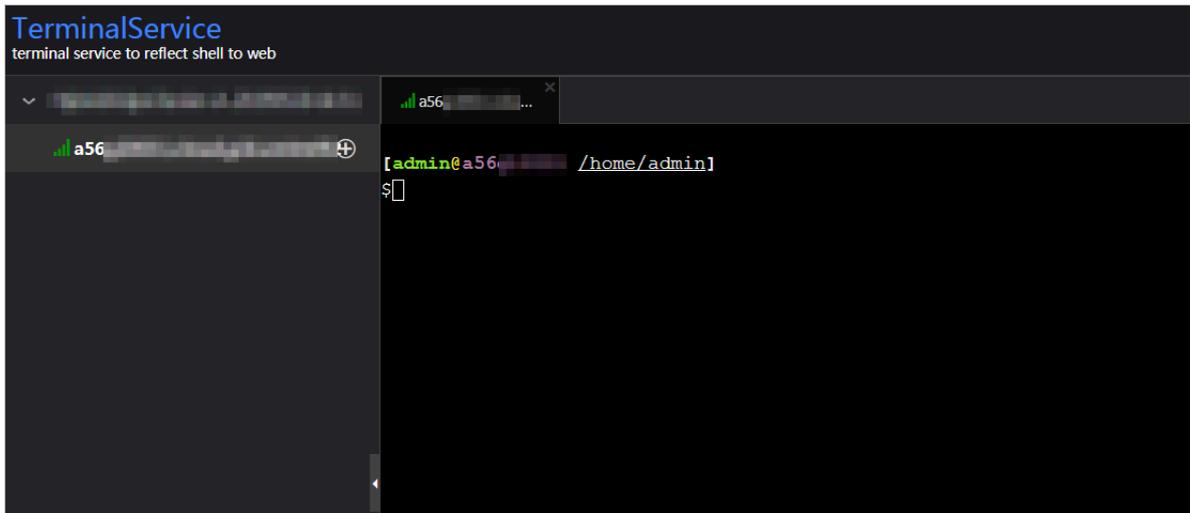
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

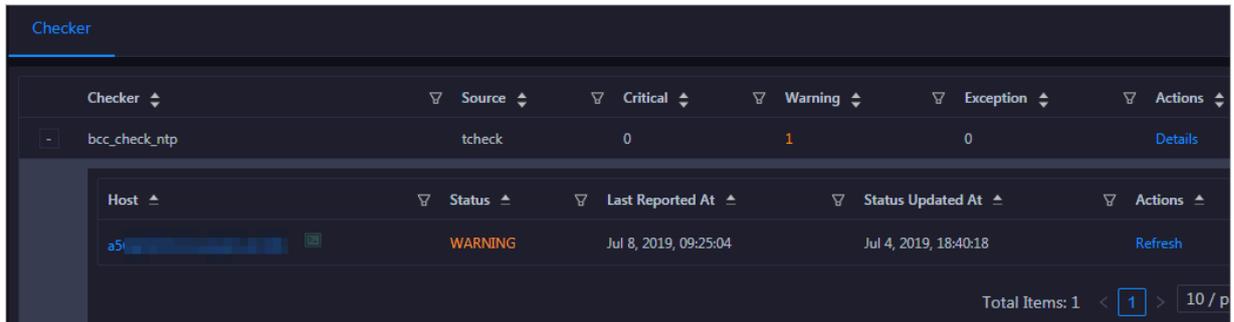


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

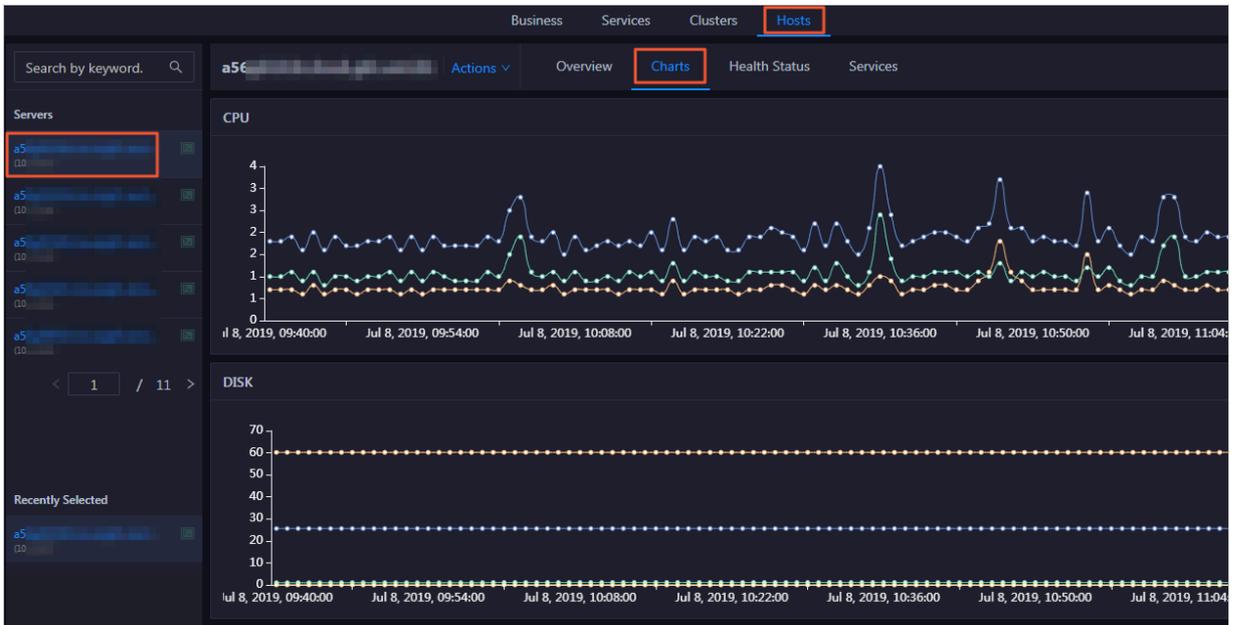
After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.6.5.3 Host charts

On the host chart page, you can view the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.

On the Hosts page, select a host in the left-side navigation pane, and then click the Charts tab. The Charts page for the host appears.



The Charts page displays trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host. For more information, see [Host overview](#).

3.1.6.5.4 Host services

On the host service page, you can view information about service instances and service instance roles of a host.

On the Hosts page, select a host in the left-side navigation pane, and then click the Services tab. The Services page for the host appears.

Cluster	ServiceInstance	Role
Blink	bigdata-sre	Agent#
Blink	blink-server	Worker#
Blink	tianji-sshtunnel-client	SSHTunnelClient#
Blink	hids-client	HidsClient#
Blink	tianji	TianjiClient#

Total Items: 5 < 1 > 10 / page Goto

On the Services page, you can view the cluster, service instances, and service instance roles of the host.

3.1.7 DataHub

3.1.7.1 O&M overview

This topic describes the features of DataHub O&M and how to access the DataHub O&M page.

Modules

DataHub O&M includes business O&M, service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature		Description
Business O&M	Projects		Displays the name, owner, the number of topics, read traffic, write traffic, storage of each project, and the time when a project is created.
	Topics		Displays the name of a topic, the name of the project to which the topic belongs, the number of shards, storage, read traffic, write traffic, and the time when the topic is created.
Service O&M	Job Scheduler	Overview	Displays the key operation metrics of Job Scheduler, including the service overview, service status, health check result, health check history, resource usage, and compute node overview. You can also view the trend charts of CPU and memory usage on this page.
		Instances	Displays information about the Job Scheduler service roles, including the name, host, IP address, and status of a service role, and host status.
		Health Status	Displays all checkers of Job Scheduler, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

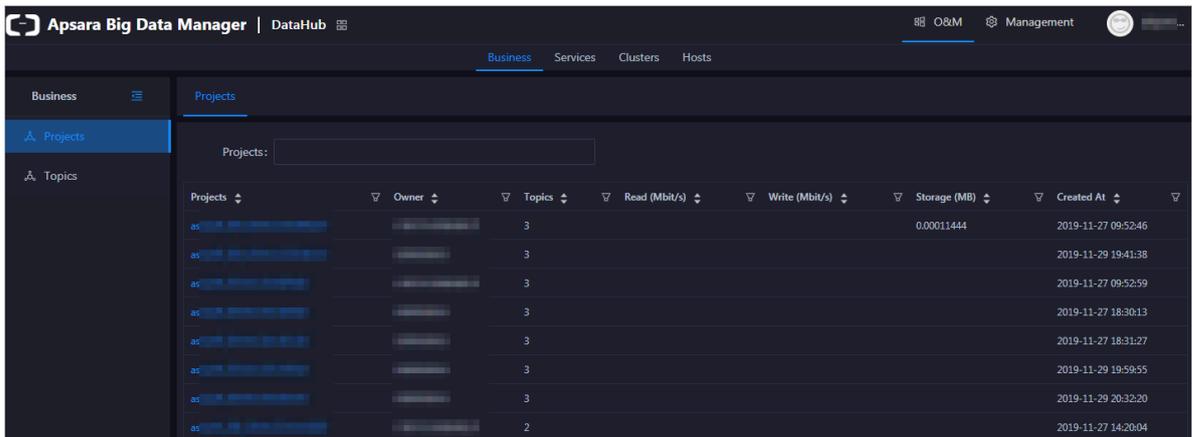
Module	Feature	Description	
	Compute Nodes	Displays all compute nodes of a cluster , including total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active. You can also add compute nodes to or remove them from the blacklist or read-only list on the Compute Nodes page.	
	Apsara Distributed File System	Overview	Displays the key operation metrics of Apsara Distributed File System, including the service overview, service status, health check result, health check history, storage usage, and storage nodes. You can also view the trend charts of storage usage and file count on this page.
		Instances	Displays information about the Apsara Distributed File System service roles , including the name, host, IP address , and status of a service role, and host status.
		Health Status	Displays all checkers of Apsara Distributed File System, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.
		Storage Nodes	Displays information about the storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, time to live (TTL), and send buffer size. You can also set the status of storage nodes and data disks on this page.

Module	Feature	Description
Cluster O&M	Overview	Displays the overall running information about a cluster, including the host status, service status, health check result, and health check history. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.
	Health Status	Displays the check results of all checkers in a cluster. The check results are divided into the Critical, Warning, Exception, and OK types.
	Hosts	Displays information about all hosts in a cluster, including the CPU usage, memory usage, root disk usage, packet loss rate, and packet error rate.
	Scale in Cluster and Scale out Cluster actions	Allow you to scale in or out a DataHub cluster by removing or adding physical hosts.
	Delete Topic from Smoke Testing action	Allows you to delete topics from a DataHub test project and view the execution history.
	Reverse Parse Request ID action	Allows you to reverse parse RequestId to obtain the time when a job was run and the IP address of the host. You can use the obtained information to query logs for troubleshooting.
Host O&M	Overview	Displays the overall running information about a host in a DataHub cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Charts	Displays the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.

Module	Feature	Description
	Health Status	Displays the checkers of all hosts, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.
	Services	Displays information about service instances and service instance roles of a host.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click DataHub.
3. On the DataHub page, click O&M in the upper-right corner. The Business page appears.



The O&M page includes four modules, namely, Business, Services, Clusters, and Hosts.

3.1.7.2 Business O&M

3.1.7.2.1 Business O&M overview

This topic describes the features of DataHub business O&M and how to access the business O&M page.

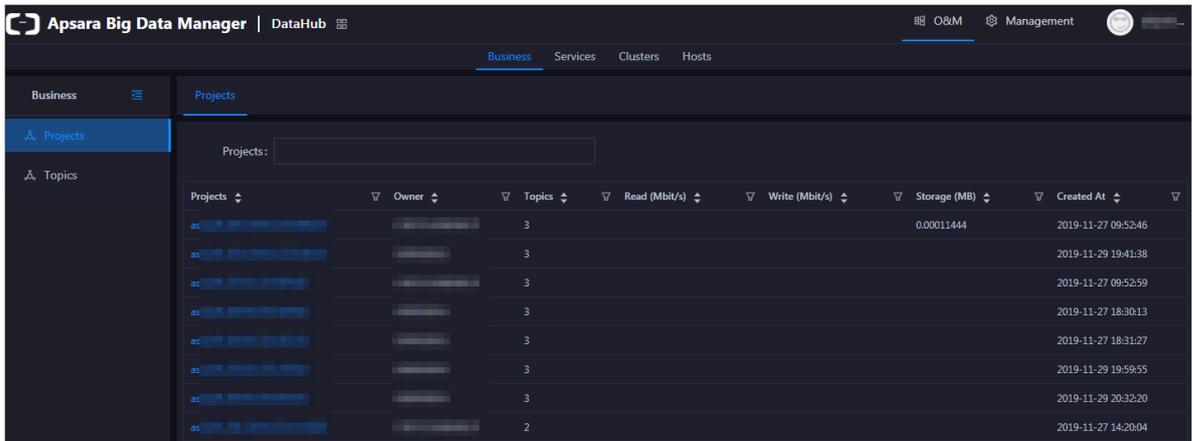
Business O&M features

- **Projects page:** displays the name, owner, the number of topics, read traffic, write traffic, storage of each project, and the time when a project is created.

- **Topics page:** displays the name of a topic, the name of the project to which the topic belongs, the number of shards, storage, read traffic, write traffic, and the time when the topic is created.

Entry

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click DataHub.
3. On the DataHub page, click O&M in the upper-right corner, and then click the Business tab. The Projects page appears.



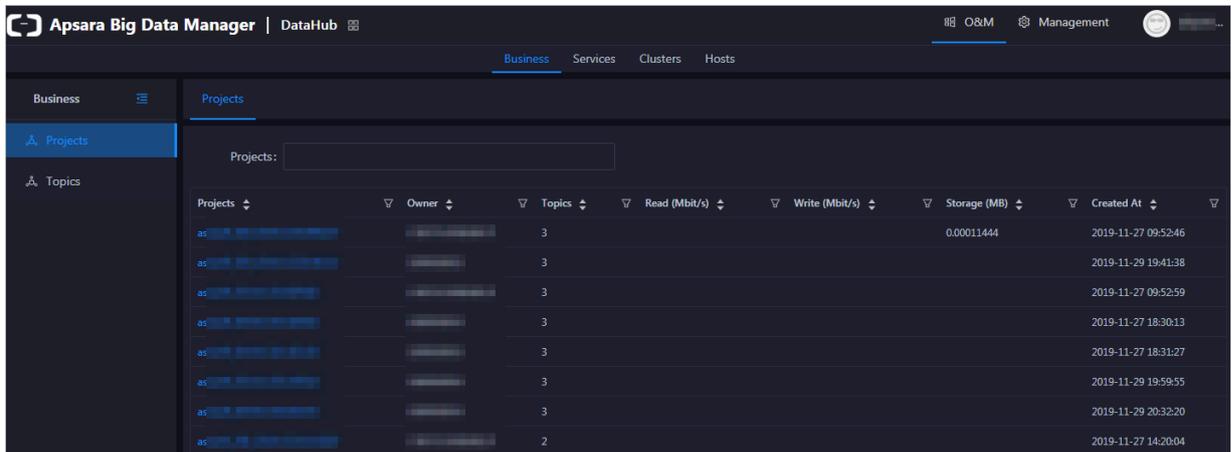
Projects	Owner	Topics	Read (Mbit/s)	Write (Mbit/s)	Storage (MB)	Created At
as-...	...	3			0.00011444	2019-11-27 09:52:46
as-...	...	3				2019-11-29 19:41:38
as-...	...	3				2019-11-27 09:52:59
as-...	...	3				2019-11-27 18:30:13
as-...	...	3				2019-11-27 18:31:27
as-...	...	3				2019-11-29 19:59:55
as-...	...	3				2019-11-29 20:32:20
as-...	...	2				2019-11-27 14:20:04

3.1.7.2.2 Projects

The Projects page displays the name, owner, the number of topics, read traffic, write traffic, storage of each project, and the time when a project is created.

Entry

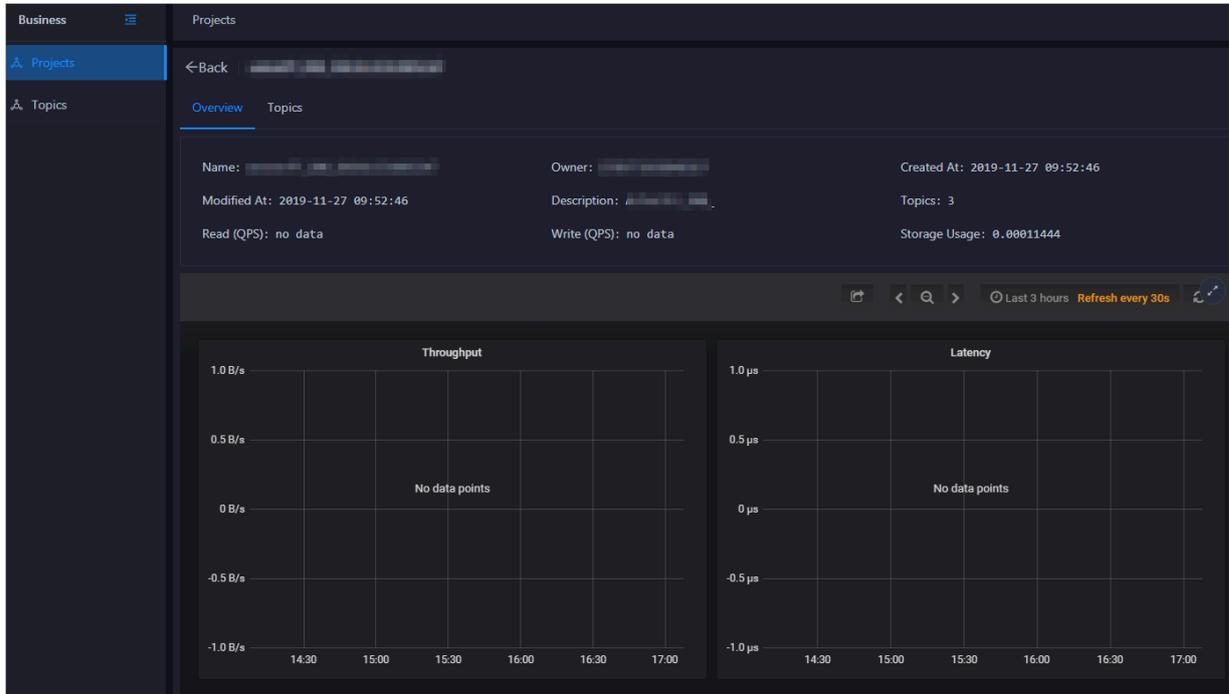
On the Business page, click Projects in the left-side navigation pane. The Projects page appears on the right.



Projects	Owner	Topics	Read (Mbit/s)	Write (Mbit/s)	Storage (MB)	Created At
as-...	...	3			0.00011444	2019-11-27 09:52:46
as-...	...	3				2019-11-29 19:41:38
as-...	...	3				2019-11-27 09:52:59
as-...	...	3				2019-11-27 18:30:13
as-...	...	3				2019-11-27 18:31:27
as-...	...	3				2019-11-29 19:59:55
as-...	...	3				2019-11-29 20:32:20
as-...	...	2				2019-11-27 14:20:04

View project overview

On the Projects page, click the name of a project that you want to view. The Overview page for the project appears.



View topics of a project

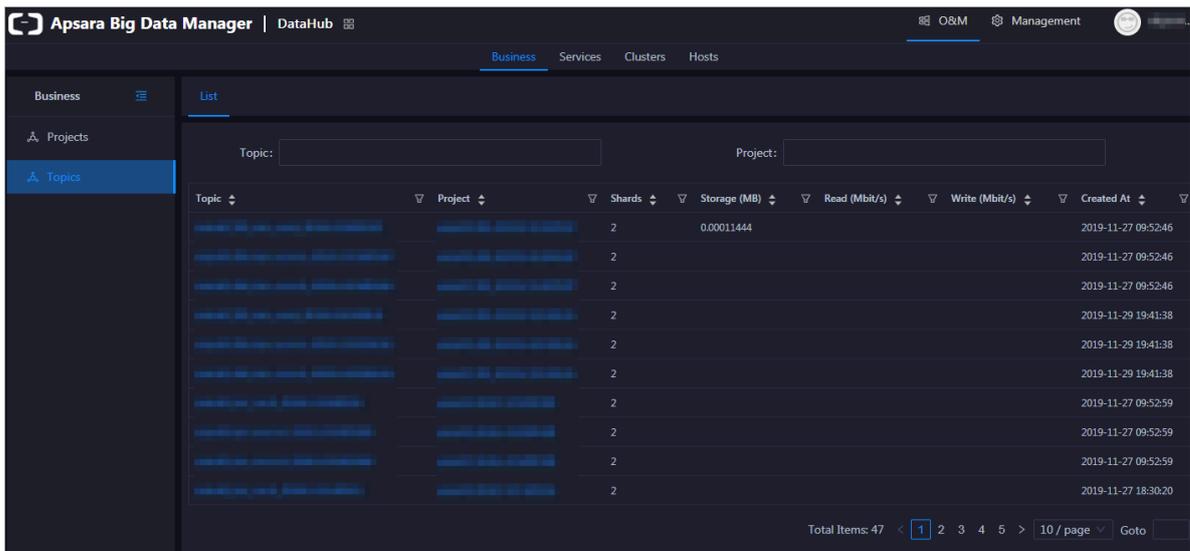
On the Projects page, click the name of a project that you want to view. On the page that appears, click the Topics tab. All topics in the project appear.

3.1.7.2.3 Topics

The Topics page displays the name of a topic, the name of the project to which the topic belongs, the number of shards, storage, read traffic, write traffic, and the time when the topic is created.

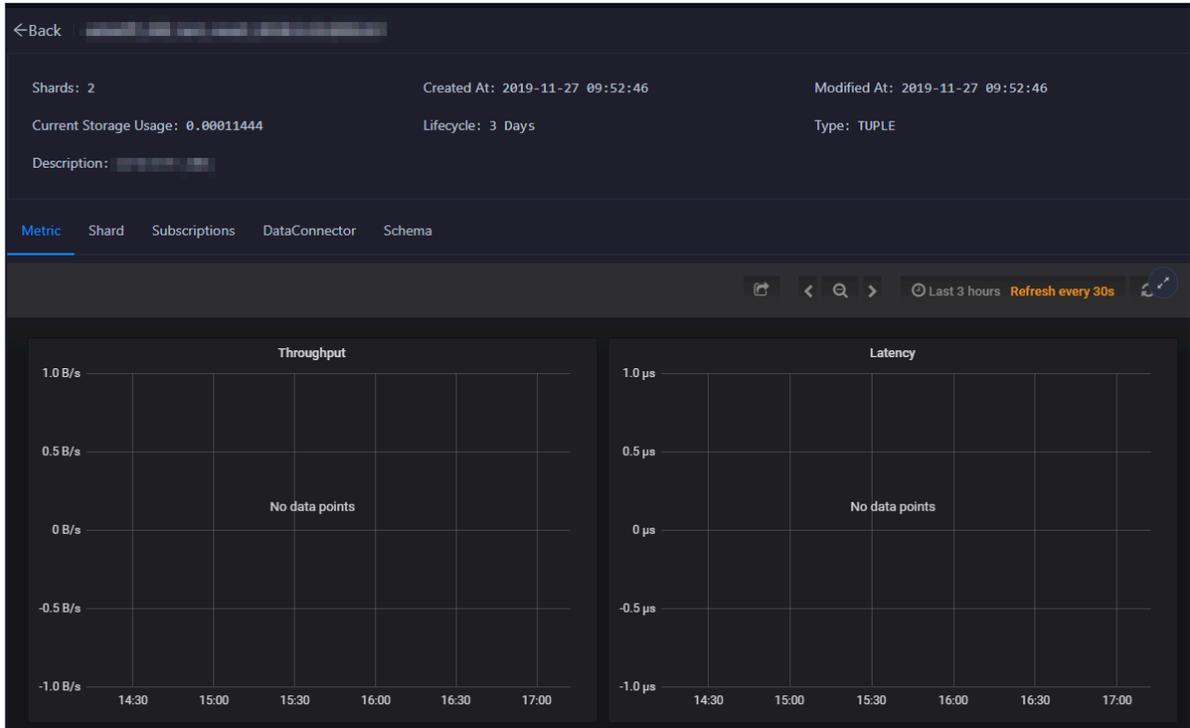
Entry

1. On the Business page, click Topics in the left-side navigation pane. The List page appears.



2. On the List page, click the name of the topic that you want to view. On the page that appears, you can view the number of shards, the time when the topic is created and modified, current storage usage, lifecycle, type, and description.

You can also view more details about monitoring metrics, shards, subscriptions, DataConnectors, and schema.



Features:

- **Metric:** You can view information about the throughput and latency of a topic in quasi-real time.
- **Shard:** Shards are concurrent tunnels used for data transmission in a topic.

You can view the ID, status, and active time of each shard.

- **Subscriptions:** The subscription feature of DataHub supports saving checkpoints to the server and restoring data from any checkpoint you saved.

You can view the ID, status, owner, and description of each subscription and the time when the subscription is modified.

- **DataConnector:** DataConnectors synchronize the streaming data from DataHub to other cloud products. You can configure a DataConnector so that the data you write to DataHub can be used in other cloud products.

You can view the name, ID, owner, and status of each DataConnector, and the time when the DataConnector is created and modified.

- **Schema:** The schema is used to define the data types of fields.

You can view the data type and name of each field.

3.1.7.3 Service O&M

3.1.7.3.1 Service O&M for Job Scheduler

3.1.7.3.1.1 Service O&M overview

This topic describes the service O&M features of Job Scheduler and how to access the service O&M page.

Modules

- **Overview page:** displays the key operation metrics of Job Scheduler, including the service overview, service status, health check result, health check history, resource usage, and compute node overview. You can also view the trend charts of CPU and memory usage on this page.
- **Instances page:** displays information about the service roles of Job Scheduler.
- **Health Status page:** displays the check results of all checkers for Job Scheduler. The check results are divided into the Critical, Warning, Exception, and OK types.
- **Compute Nodes page:** displays all Job Scheduler compute nodes and allows you to add compute nodes to or remove them from the blacklist or read-only list.

3.1.7.3.1.2 Service overview

The Overview page displays the key operation metrics of Job Scheduler, including the service overview, service status, health check result, health check history, resource usage, and compute node overview. You can also view the trend charts of CPU and memory usage on this page.

Entry

1. On the Services page, click **Fuxi** in the left-side navigation pane.
2. Select a cluster from the drop-down list, and then click the **Overview** tab. The Overview page for Job Scheduler appears.

The Overview page displays the key operation metrics of Job Scheduler, including the service overview, service status, health check result, health check history, resource usage, and compute node overview. You can also view the trend charts of CPU and memory usage on this page.

Services

This section displays the numbers of available services, unavailable services, and services that are being upgraded respectively.

Roles

This section displays all Job Scheduler service roles and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

Saturability - Resource Usage

This section displays the usage and allocation of CPU and memory resources.

- **CPU (Core):** displays the CPU usage, the total number of CPU cores, the number of available CPU cores, and the CPU cores for SQL acceleration.
- **Memory (Bytes):** displays the memory usage, the total memory size, the available memory size, and the size of memory for SQL acceleration.

CPU Usage (1/100 Core) and Memory Usage (MB)

This section displays trend charts of CPU and memory usage for Job Scheduler . Each trend chart displays the trend lines of the used quota, minimum quota , maximum cluster quota, requested quota, and maximum quota over time in different colors.

Click  in the upper-right corner of a chart to zoom in it. The following figure shows an enlarged CPU usage chart.

You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.

Compute Nodes

This section displays the details of Job Scheduler compute nodes, including the online rate, the total number of compute nodes, the number of online compute nodes, and the number of compute nodes in the blacklist.

Health Check

This section displays the number of checkers deployed for Job Scheduler and the respective number of Critical, Warning, and Exception alerts.

Click **View Details** to go to the **Health Status** page. On this page, you can view the health check details. For more information, see [Cluster health](#).

Health Check History

This section displays a record of the health checks performed on Job Scheduler. You can view the respective number of Critical, Warning, and Exception alerts for each health check.

Click **View Details** to go to the **Health Status** page. On this page, you can view the health check details. For more information, see [Cluster health](#).

You can click the event content of a check to view the exception items.

3.1.7.3.1.3 Service instances

The **Instances** page displays information about the Job Scheduler service roles, including the name, host, IP address, and status of a service role, and host status.

1. On the **Services** page, click **Fuxi** in the left-side navigation pane.
2. Select a cluster from the drop-down list, and then click the **Instances** tab. The **Instances** page for Job Scheduler appears.

On the **Instances** page, you can view information about the Job Scheduler service roles, including the name, host, IP address, and status of a service role, and host status.

3.1.7.3.1.4 Health status

On the **Health Status** page for Job Scheduler, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. On the **Services** page, click **Fuxi** in the left-side navigation pane.
2. Select a cluster from the drop-down list, and then click the **Health Status** tab. The **Health Status** page for Job Scheduler appears.

On the **Health Status** page, you can view all checkers of the Job Scheduler service and the check results for all hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are

alerts. You need to pay attention to them, especially the Critical and Warning results.

Other operations

On the Health Status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host. For more information, see [Cluster health](#).

3.1.7.3.1.5 Compute nodes

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active. You can also add compute nodes to or remove them from the blacklist or read-only list on the Compute Nodes page.

Entry

- 1. On the Services page, click Fuxi in the left-side navigation pane.**
- 2. Select a cluster from the drop-down list, and then click the Compute Nodes tab. The Compute Nodes page for Job Scheduler appears.**

You can view the details of compute nodes on the Compute Nodes page for Job Scheduler, including total CPU, idle CPU, total memory, and idle memory of each compute node. You can also check whether a node is added to the blacklist and whether it is active.

Blacklist and read-only setting

You can add compute nodes to or remove them from the blacklist or read-only list. To add compute nodes to the blacklist, follow these steps:

- 1. On the Compute Nodes page, click Actions for the target compute node and then select Add to Blacklist.**
- 2. In the dialog box that appears, enter the hostname. If you want to add multiple compute nodes to the blacklist, separate hostnames with commas (,).**
- 3. Click Run. A message appears, indicating that the action has been submitted.**

You can view the blacklist statuses of compute nodes in the compute node list after the configuration is completed.

3.1.7.3.2 Service O&M for Apsara Distributed File System

3.1.7.3.2.1 Service O&M overview

This topic describes the service O&M features of Apsara Distributed File System and how to access the service O&M page.

Modules

- **Overview page:** displays the key operation metrics of Apsara Distributed File System, including the service overview, service status, health check result, health check history, storage usage, and storage nodes. You can also view the trend charts of storage usage and file count on this page.
- **Instances page:** displays information about the service roles of Apsara Distributed File System.
- **Health Status page:** displays the check results of all checkers for Apsara Distributed File System. The check results are divided into the Critical, Warning, Exception, and OK types.
- **Storage Nodes page:** displays information about the storage nodes of Apsara Distributed File System. On this page, you can set the status of a storage node to Disabled or Normal. In addition, you can set the status of a disk on a storage node to Normal or Error.
- **Change Primary Master Node action:** allows you to change the primary master node of Apsara Distributed File System in a cluster.
- **Empty Recycle Bin action:** allows you to empty the recycle bin of Apsara Distributed File System.
- **Enable Data Rebalancing or Disable Data Rebalancing action:** allows you to enable or disable the data rebalancing feature of Apsara Distributed File System.
- **Run Checkpoint on Master Node action:** allows you to run checkpoints on master nodes of Apsara Distributed File System to write memory data into disks.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click DataHub.
3. On the DataHub page that appears, click O&M in the upper-right corner, and then click the Services tab.

4. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster from the drop-down list. The Overview page for Apsara Distributed File System appears.

3.1.7.3.2 Service overview

On the Overview page for Apsara Distributed File System, you can view the key operation metrics, including the service overview, service status, storage usage, and storage node overview. You can also view the trend charts of storage usage and file count on this page.

Entry

1. On the Services page, click Pangu in the left-side navigation pane.
2. Select a cluster from the drop-down list, and then click the Overview tab. The Overview page for Apsara Distributed File System appears.

The Overview page displays the key operation metrics of Apsara Distributed File System, including the service overview, service status, health check result, health check history, storage usage, and storage nodes. You can also view the trend charts of storage usage and file count on this page.

Services

This section displays the status of Apsara Distributed File System and the number of service roles.

Roles

This section displays all Apsara Distributed File System service roles and their statuses. You can also view the expected and actual numbers of machines in the final status for each service role.

Saturability - Storage

This section displays the storage usage (Storage) and file count (File Count).

- **Storage:** displays the storage usage, total storage size, available storage size, and recycle bin size.
- **File Count:** displays the file count usage, maximum number of files, number of existing files, and number of files in the recycle bin.

Storage Nodes

This section displays information about the storage nodes of Apsara Distributed File System, including the respective number of data nodes, normal nodes, disks, and normal disks. You can also view the faulty node percentage and faulty disk percentage in this section.

Health Check

This section displays the number of checkers deployed for Job Scheduler and the respective number of Critical, Warning, and Exception alerts.

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

Health Check History

This section displays a record of the health checks performed on Job Scheduler. You can view the respective number of Critical, Warning, and Exception alerts for each health check.

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

You can click the event content of a check to view the exception items.

3.1.7.3.2.3 Service instances

The Instances page displays information about the Apsara Distributed File System service roles, including the name, host, IP address, and status of a service role, and host status.

Entry

- 1. On the Services page, click Pangu in the left-side navigation pane.**
- 2. Select a cluster from the drop-down list, and then click the Instances tab. The Instances page for Apsara Distributed File System appears.**

On the Instances page, you can view information about the Apsara Distributed File System service roles, including the name, host, IP address, and status of a service role, and host status.

Other operations

You can filter or sort service roles by column to facilitate information display. For more information, see [Common operations](#).

3.1.7.3.2.4 Health status

On the Health Status page for Apsara Distributed File System, you can view all health check items, check details, check results, and solutions for alerts (if there are any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. On the Services page, click Pangu in the left-side navigation pane.
2. Select a cluster, and then click the Health Status tab. The Health Status page for Apsara Distributed File System appears.

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

Other operations

On the Health Status page, you can view checker details, hosts with alerts, and alert causes. You can also log on to hosts with alerts, clear alerts, and run checkers again. For more information, see [Cluster health](#).

3.1.7.3.2.5 Storage nodes

The Storage Nodes page displays the information about the storage nodes of Apsara Distributed File System, including the total storage size, available storage size, status, time to live (TTL), and send buffer size. You can also set the status of storage nodes and data disks on this page.

Entry

1. On the Services page, click Pangu in the left-side navigation pane.
2. Select a cluster, and then click the Storage Nodes tab. The Storage Nodes page for Apsara Distributed File System appears.

On the Storage Nodes page, you can view storage node details, including the total CPU, available CPU, total storage size, and available storage size. You can also check whether a node is added to the blacklist and whether it is active.

Set the storage node status

You can set the storage node status to Disabled or Normal. This section describes how to set the status of a storage node to Disabled.

- 1. On the Storage Nodes page, click Actions for the target storage node, and then select Set Node Status to Disabled.**
- 2. In the dialog box that appears, set volume and hostname. If you need to add multiple storage nodes to the blacklist, separate the storage node names with commas (,).**
- 3. Click Run. A message appears, indicating that the action has been submitted.**

You can view the status of storage nodes in the storage node list.

Set the data disk status

You can set the data disk status to Error or Normal. This section describes how to set the status of a data disk to Error.

- 1. On the Storage Nodes page, click Actions for the target data disk, and then select Set Disk Status to Error.**
- 2. In the dialog box that appears, set volume, hostname, and diskid. If you need to add multiple data disks to the blacklist, separate the disk names with commas (,).**
- 3. Click Run. A message appears, indicating that the action has been submitted.**

3.1.7.3.2.6 Empty the recycle bin of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to empty the recycle bin of Apsara Distributed File System.

Prerequisites

Your Apsara Bigdata Manager (ABM) account has the permission to manage DataHub.

Procedure

- 1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster from the drop-down list. The Overview page for Apsara Distributed File System appears.**
- 2. Click Actions in the upper-left corner, and then click Empty Recycle Bin.**
- 3. In the dialog box that appears, set volume. The default value is PanguDefaultVolume.**

4. Click Run. A message appears, indicating that the action has been submitted.
5. View the execution status.

Click Actions in the upper-left corner, and then click  next to Empty Recycle Bin to view the execution history.

In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

6. Click Details for a failed execution to locate the failure cause.

You can also view information about parameter settings, hosts, script, and execution parameters to locate the failure cause.

3.1.7.3.2.7 Enable or disable data rebalancing for Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to enable or disable data rebalancing for Apsara Distributed File System.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on DataHub.

Disable data rebalancing

1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.
2. Click Actions, and then select Disable Data Rebalancing.
3. In the dialog box that appears, set volume. The default value is PanguDefaultVolume.
4. Click OK. A message is displayed, indicating that the action has been submitted.
5. View the execution status.

Click Actions, and then click  next to Disable Data Rebalancing to view the execution history.

In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

6. Click **Details** for a failed execution to locate the failure cause. For more information, see [Locate the failure cause](#).

Enable data rebalancing

1. On the **Services** page, click **Pangu** in the left-side navigation pane, and then select a cluster. The **Overview** page for **Apsara Distributed File System** appears.
2. Click **Actions**, and then select **Enable Data Rebalancing**.
3. In the dialog box that appears, set volume. The default value is **PanguDefaultVolume**.
4. Click **OK**. A message is displayed, indicating that the action has been submitted.
5. View the execution status.

Click **Actions**, and then click  next to **Enable Data Rebalancing** to view the execution history.

In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

6. Click **Details** for a failed execution to locate the failure cause. For more information, see [Locate the failure cause](#).

Locate the failure cause

The following describe how to locate the failure cause for enabling data rebalancing.

1. In the **Enable Data Rebalancing Execution History** dialog box, click **Details** a failed execution.
2. On the page that appears, click **View Details** for a failed step to locate the failure cause.

You can also view information about parameter settings, hosts, script, and execution parameters to locate the failure cause.

3.1.7.3.2.8 Run a checkpoint on the master nodes of Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to run checkpoints on master nodes of Apsara Distributed File System. This operation writes memory data into disks.

When a failure occurs to Apsara Distributed File System, you can use checkpoints to restore data to the status before the failure. This ensures data consistency.

Prerequisites

Your ABM account must have the required permissions to perform O&M operations on DataHub.

Procedure

1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.
2. Click Actions, and then select Run Checkpoint on Master Node.
3. In the dialog box that appears, set volume. The default value is PanguDefaultVolume.
4. Click OK. A message is displayed, indicating that the action has been submitted.
5. View the execution status.

Click Actions, and then click  next to Run Checkpoint on Master Node to view the execution history.

In the Current Status column, RUNNING indicates that the execution is in progress, SUCCESS indicates that the execution is successful, and FAILED indicates that the execution fails.

6. Click Details for a failed execution to locate the failure cause.

You can also view information about the parameter settings, hosts, script, and execution parameters to locate the failure cause.

3.1.7.3.2.9 Change the primary master node for Apsara Distributed File System

Apsara Bigdata Manager (ABM) allows you to perform primary/secondary switchover on the master nodes of Apsara Distributed File System. After the primary/secondary switchover is completed, a secondary master node becomes the new primary master node, and the original primary master node becomes a new secondary master node.

Prerequisites

- Your ABM account must have the required permissions to perform O&M operations on DataHub.

- You have obtained the roles of the primary and secondary master nodes in a volume. To view the role of a master node, log on to the Apsara Infrastructure Management Framework console and access the PanguTools# host in the DataHub cluster. Then, run the `puadmin gems` command on the host.
- You have obtained the hostname of the secondary master node that is to be the new primary master node. Log on to the ABM console, go to the DataHub O&M page, and then click Services. On the page that appears, click Pangu in the left-side navigation pane, and then click the Instances tab. On the Instances page, view the hostnames of PanguMaster# hosts.

Context

A volume in Apsara Distributed File System is similar to a namespace in Hadoop Distributed File System (HDFS). The default volume is PanguDefaultVolume. Multiple volumes may exist if a cluster consists of numerous nodes. A volume has three master nodes. One of the nodes serves as the primary master node, whereas the other two nodes serve as secondary master nodes.

Procedure

1. On the Services page, click Pangu in the left-side navigation pane, and then select a cluster. The Overview page for Apsara Distributed File System appears.
2. Click Actions, and then select Change Primary Master Node.
3. In the dialog box that appears, set relevant parameters.

The parameters are described as follows:

- **Volume:** the volume whose primary master node needs to be changed. Default value: PanguDefaultVolume. If a cluster consists of multiple volumes, set this parameter to the name of the actual volume whose primary master node needs to be changed.
 - **Hostname:** the hostname of the secondary master node that is to be the new primary master node.
 - **log_gap:** the maximum log number gap between the original primary and secondary master nodes. During the switchover, the system checks the log number gap between the original primary and secondary master nodes. If the gap is less than the specified value, switchover is allowed. Otherwise, you cannot change the primary master node. Default value: 100000.
4. Click Run. A message appears, indicating that the action has been submitted.

5. View the execution status.

Click **Actions**, and then click  next to **Change Primary Master Node** to view the execution history.

In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

6. Click **Details** for a failed execution to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

3.1.7.4 Cluster O&M

3.1.7.4.1 Cluster O&M overview

This topic describes the features of DataHub cluster O&M and how to access the cluster O&M page.

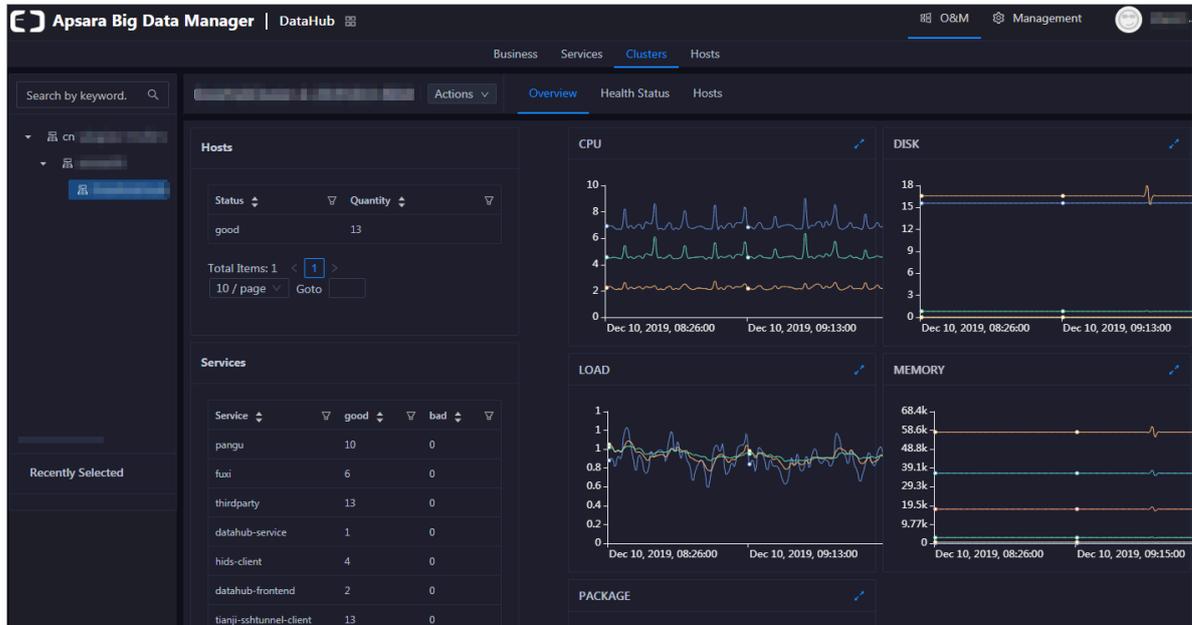
Cluster O&M features

- **Overview page:** displays the overall running information about a cluster. On this page, you can view the host status, service status, health check result, and health check history. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.
- **Health Status page:** displays the check results of all checkers in a cluster. The check results are divided into the Critical, Warning, Exception, and OK types.
- **Hosts page:** displays information about all hosts in a cluster, including the CPU usage, memory usage, root disk usage, packet loss rate, and packet error rate.
- **Scale in Cluster and Scale out Cluster actions:** allow you to scale in or out a DataHub cluster by removing or adding physical hosts.
- **Delete Topic from Smoke Testing action:** allows you to delete topics from a DataHub test project and view the execution history.
- **Reverse Parse Request ID action:** allows you to reverse parse RequestId in DataHub to obtain the time when a job was run and the IP address of the host. You can use the obtained information to query logs for troubleshooting.

Entry

1. [Log on to the ABM console.](#)

2. Click  in the upper-left corner, and then click DataHub.
3. On the DataHub page that appears, click O&M in the upper-right corner, and then click the Clusters tab.
4. On the Clusters page, select a cluster in the left-side navigation pane. The Overview page for the cluster appears.



3.1.7.4.2 Cluster overview

The cluster overview page displays the overall running and health check information about a cluster. On this page, you can view the health check result and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

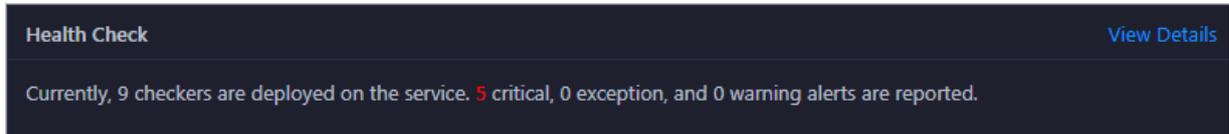
Entry

1. At the top of the O&M page, click Clusters.
2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.

The Overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the cluster. On this page, you can also view the health check result and health check history of the cluster. To view information about a cluster, select a region in the left-side navigation pane, and then select a cluster in the region.

Health Check

This section displays the number of checkers deployed for the cluster and the respective number of Critical, Warning, and Exception alerts.



Click [View Details](#) to go to the [Cluster health](#) page. On this page, you can view the health check details.

Health Check History

This section displays a record of the health checks performed on the cluster.

Click [View Details](#) to go to the [Cluster health](#) page. On this page, you can view the health check details.

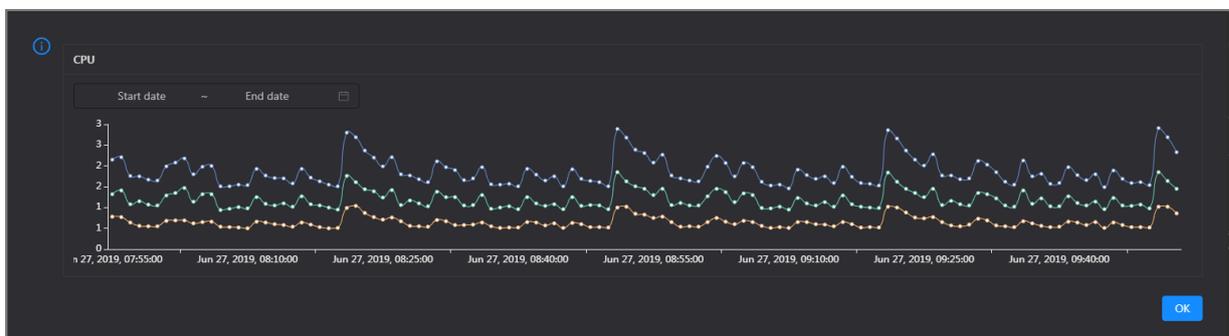
You can click the event content of a check to view the exception items.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

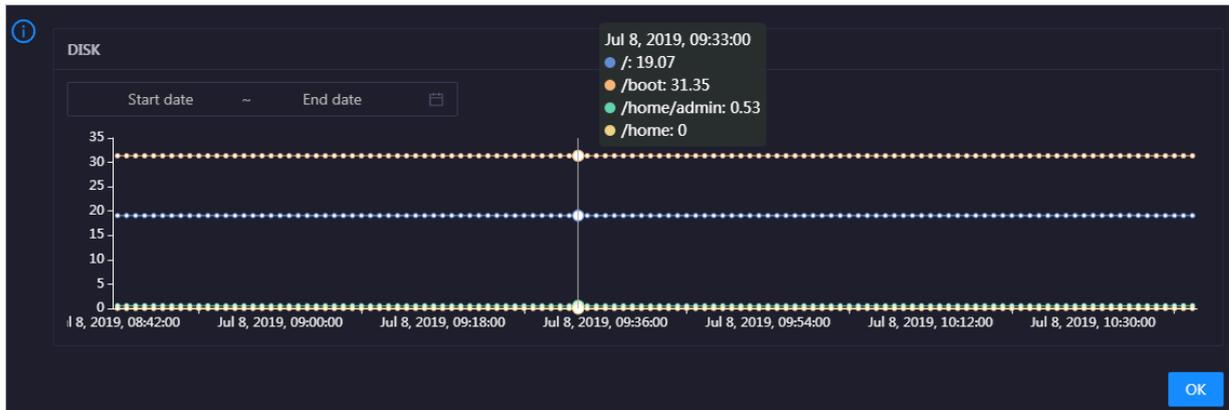
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

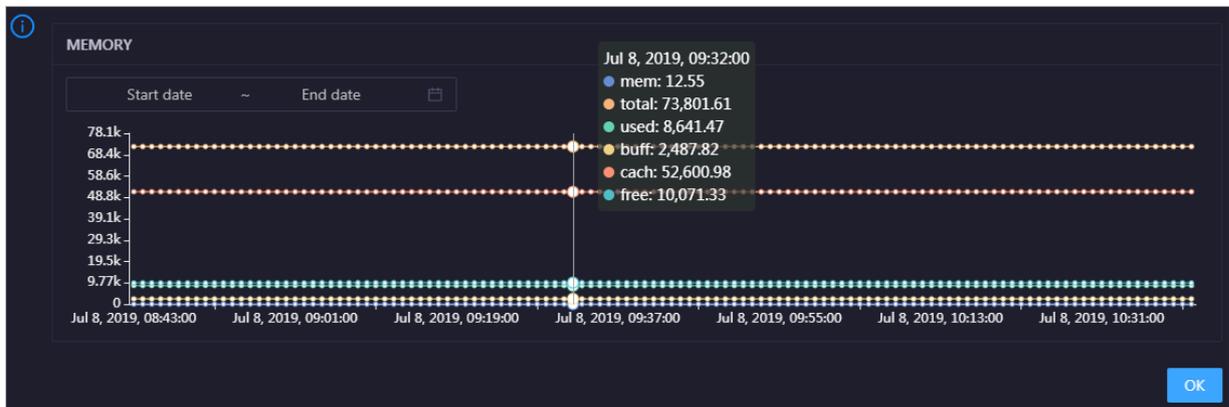


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

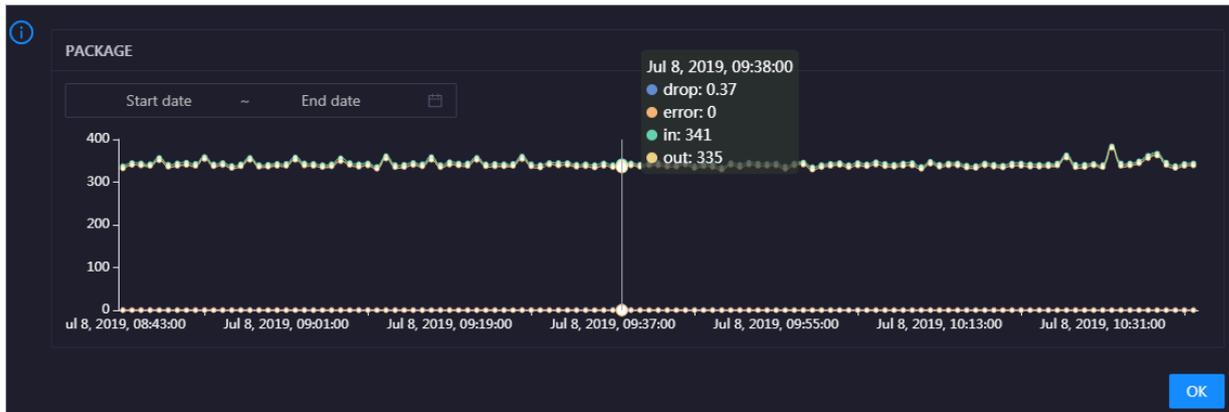


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

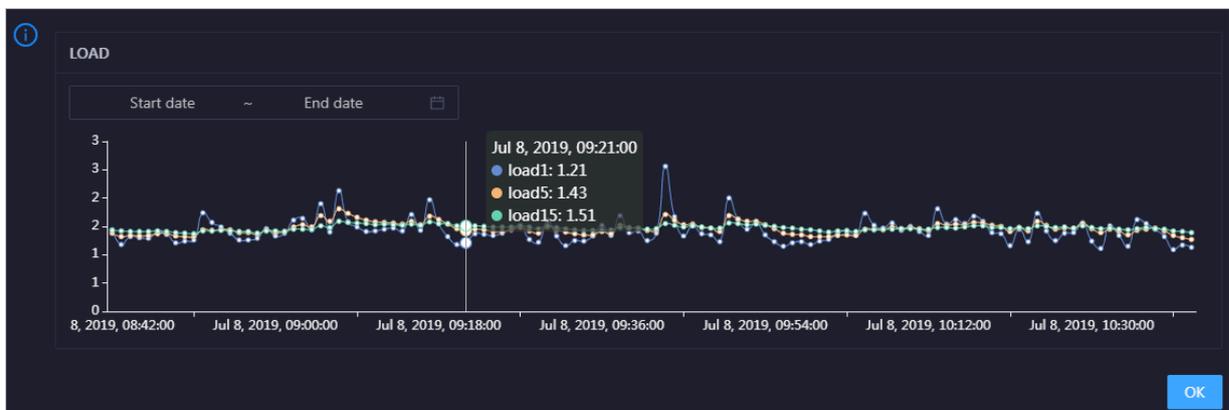


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



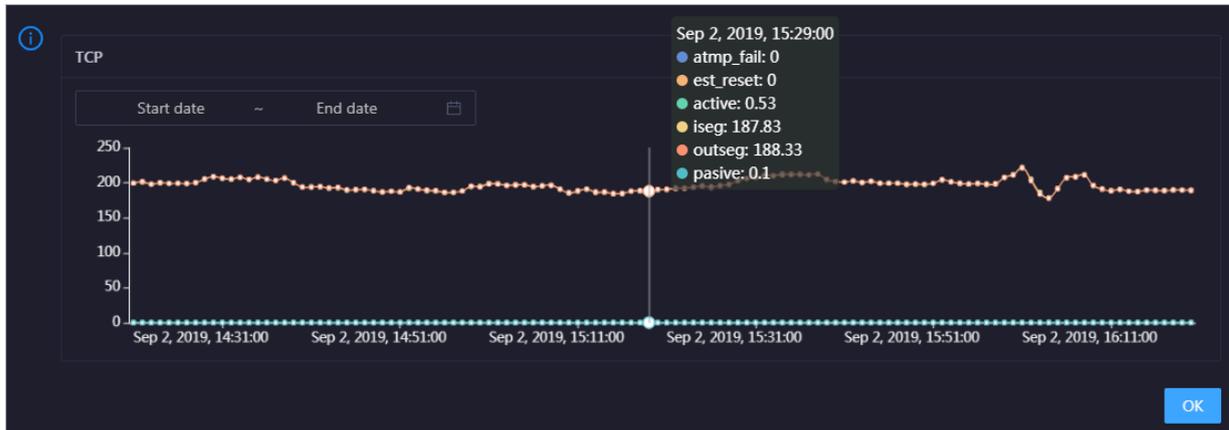
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of

sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

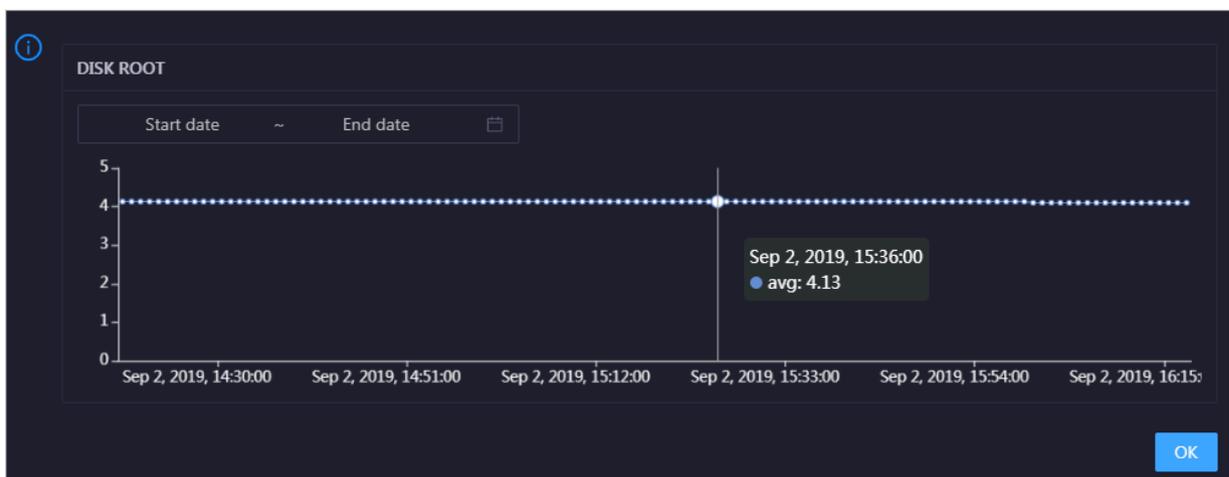


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

3.1.7.4.3 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear

alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

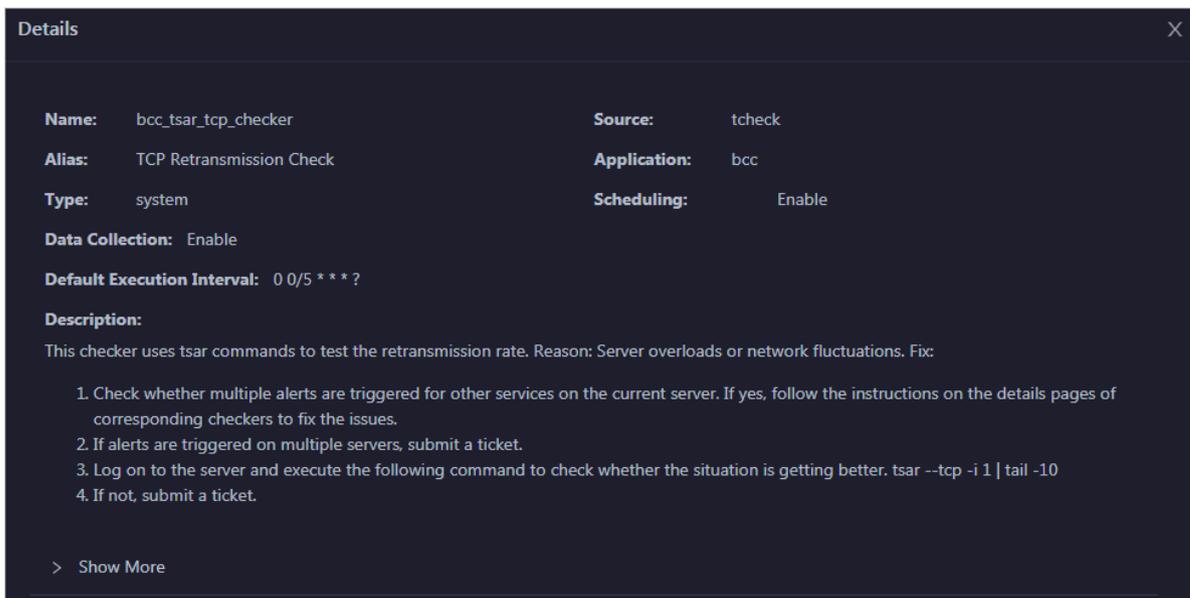
Entry

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

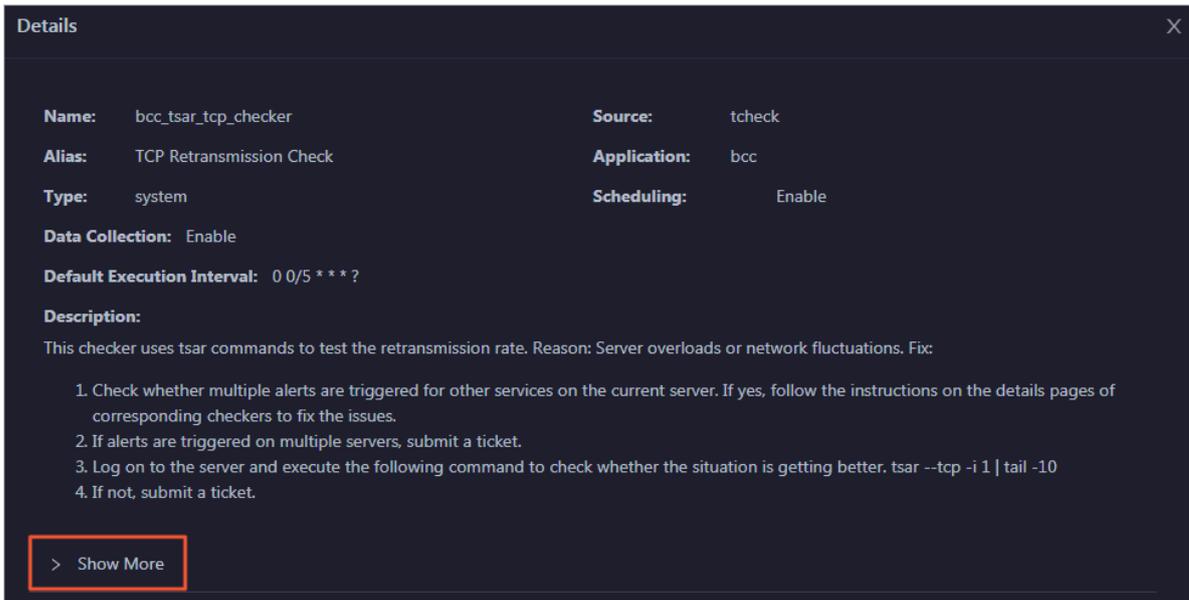
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

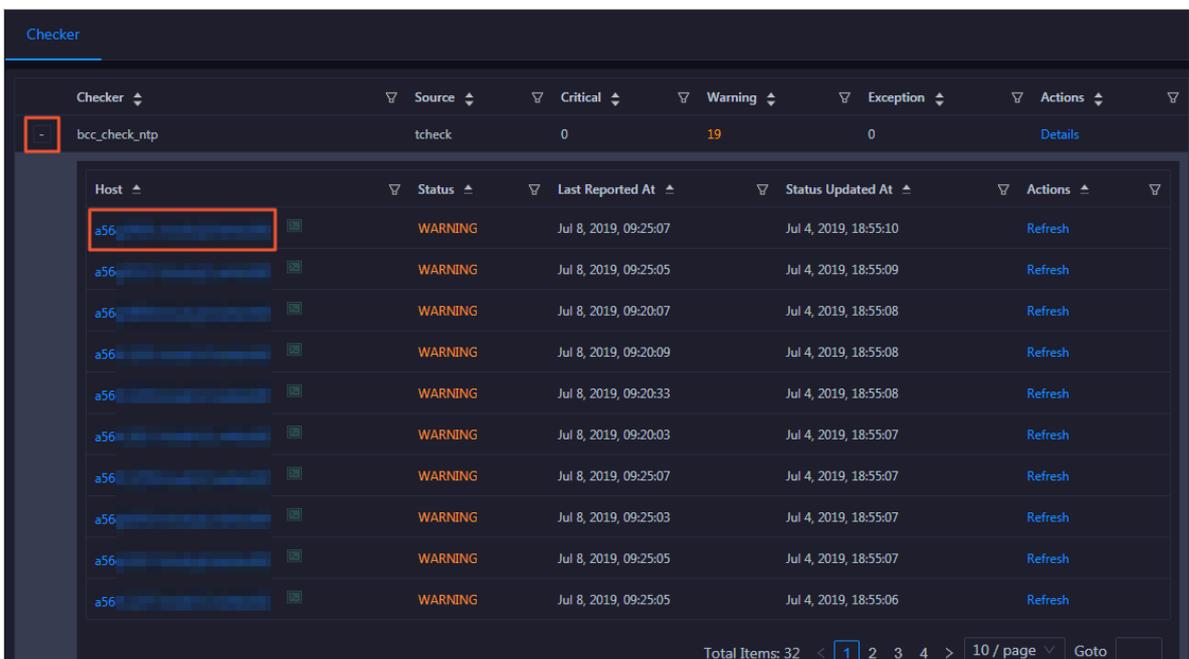


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

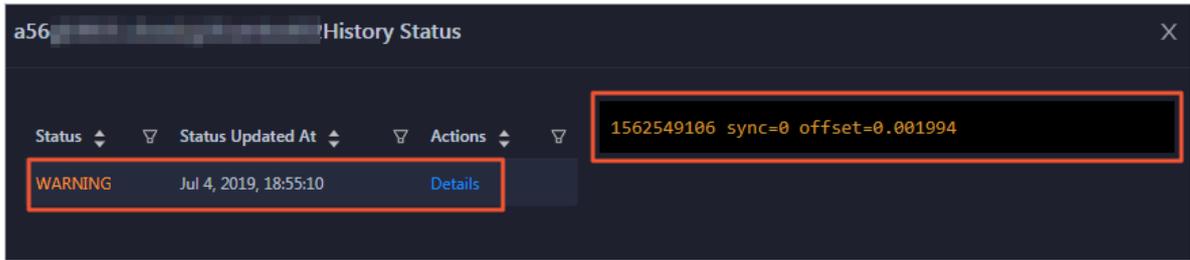
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.

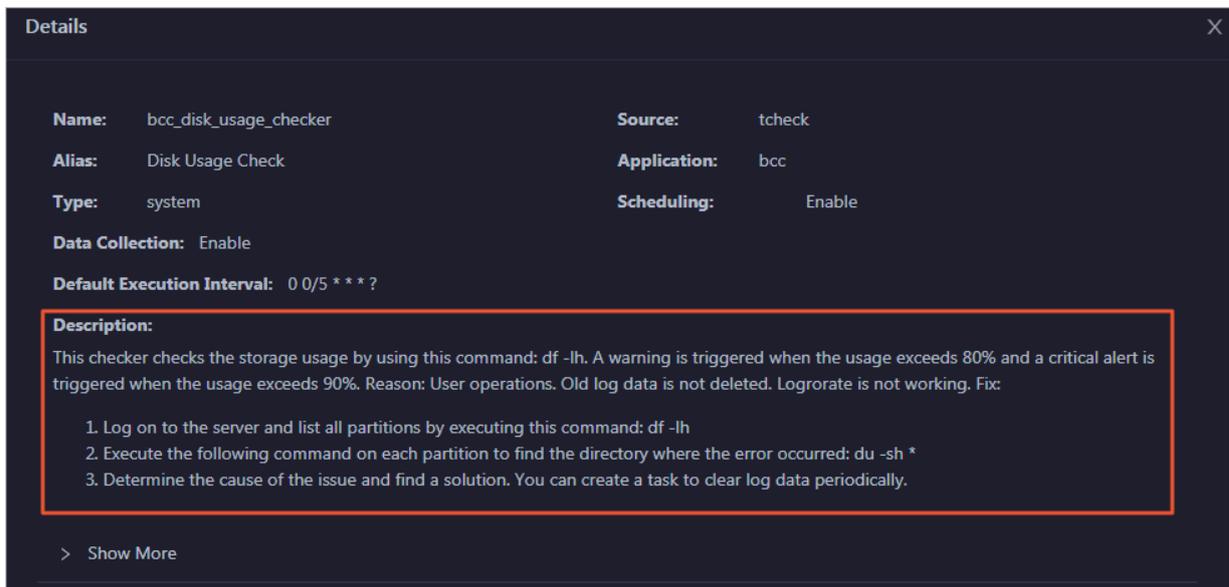


2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

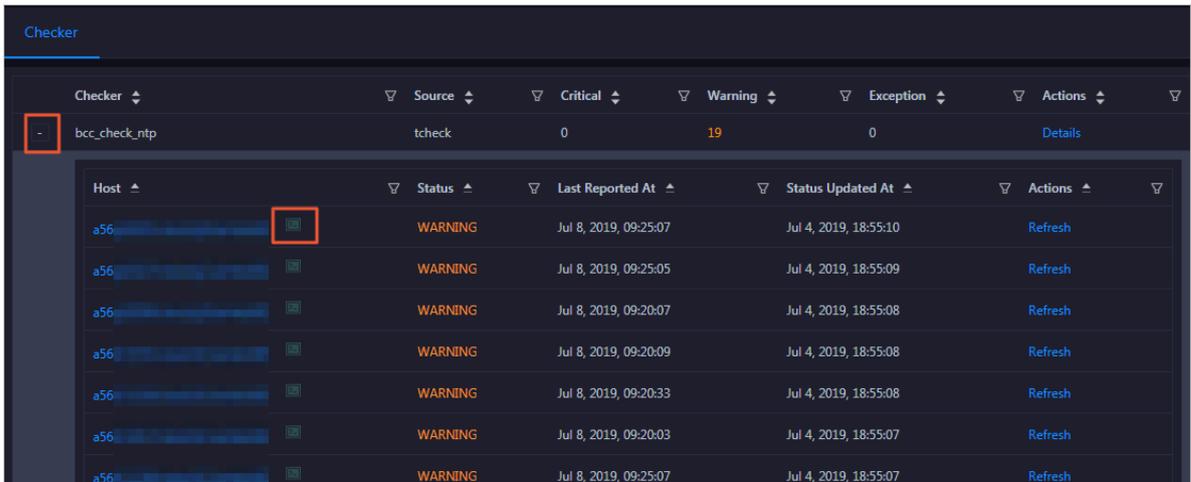
- On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



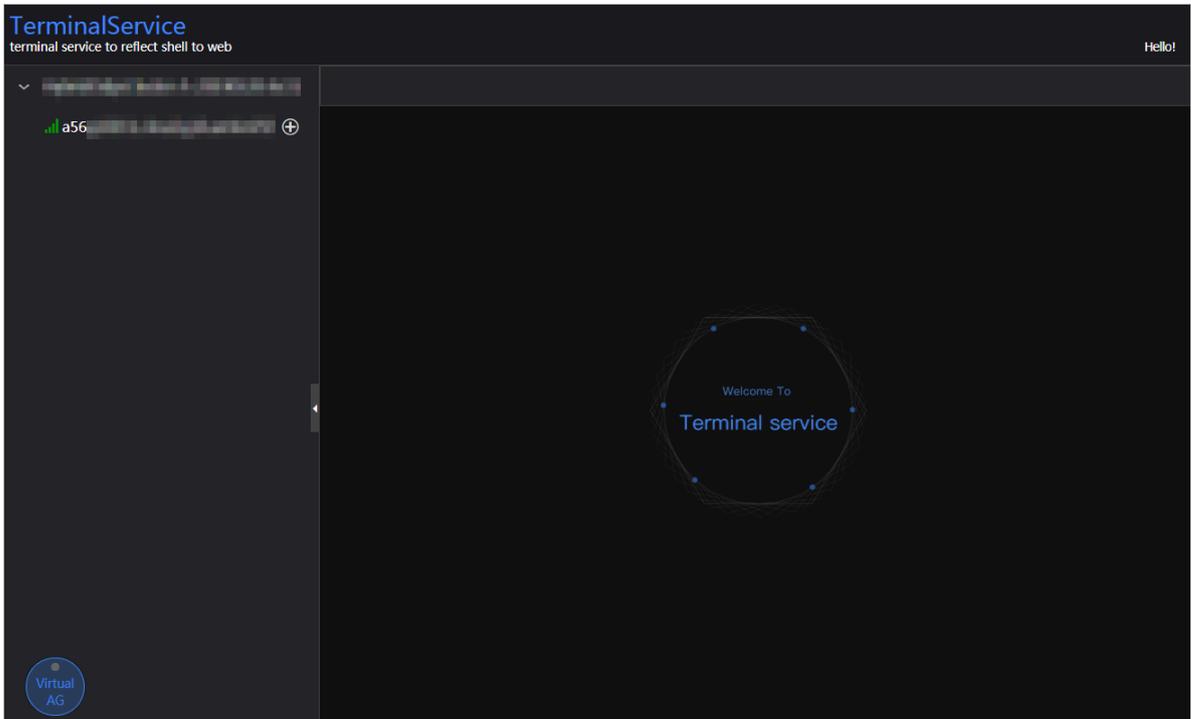
Log on to a host

- To log on to a host to clear alerts or perform other operations, follow these steps:

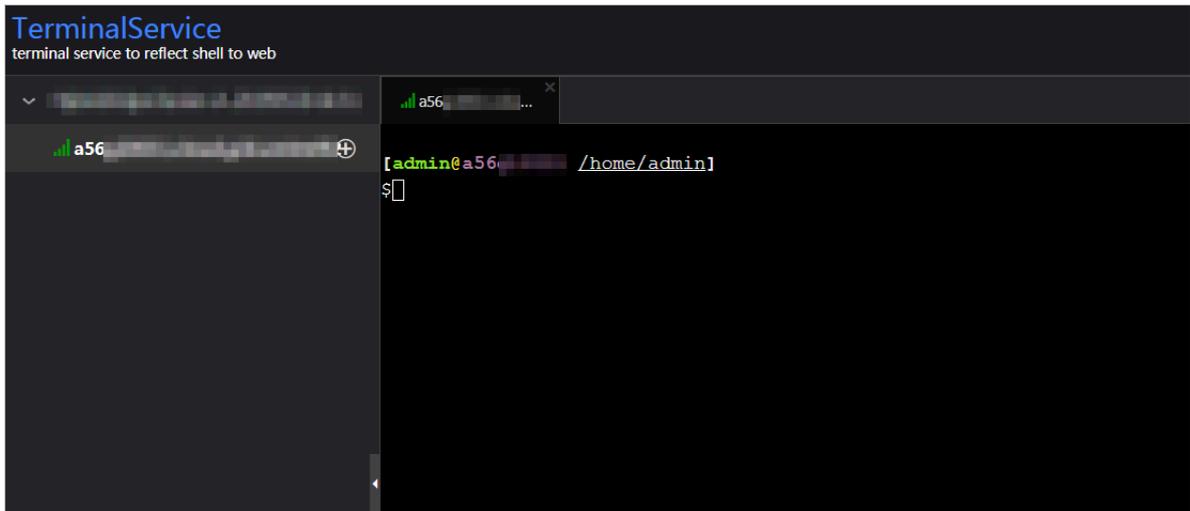
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.



3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

Checker	Source	Critical	Warning	Exception	Actions
bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

3.1.7.4.4 Cluster hosts

The cluster hosts page displays information about hosts, including the hostname, IP address, role, type, CPU usage, total memory size, available memory size, load, root disk usage, packet loss rate, and packet error rate.

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Hosts tab. The Hosts page for the cluster appears.

Hostname	IP	Role	Type	CPU Usage (%)	Total Memory (MB)	Idle Memory (MB)	Load1	Root Disk Usage (%)	Packet Loss Rate	Packet Error Rate
a56	10.	BigGraphWorker	Q41.2B	1	270685.86	225428.58	0.3	24.7	0	0
a56	10.	BigGraphWorker	Q41.2B	1.1	270685.86	222629.45	0.2	24.6	0	0
a56	10.	BigGraphWorker	Q41.2B	1	270685.86	219430.3	0.2	24.6	0	0
a56	10.	OdpsComputer	Q45.2B	1.1	115866.53	13021.39	0.7	26.5	0	0
a56	10.	OdpsComputer	Q45.2B	1.2	115866.53	14423.42	0.2	26.2	0	0
a56	10.	OdpsComputer	Q45.2B	1.3	115866.53	11324.58	0.6	26.3	0	0
a56	10.	OdpsComputer	Q45.2B	1.6	115866.53	15583.15	0.5	26.2	0	0
a56	10.	OdpsComputer	Q45.2B	1.5	115866.53	8582.05	0.5	26.5	0	0
a56	10.	OdpsComputer	Q45.2B	1.5	115866.53	14608.04	1	26.4	10	0
a56	10.	OdpsComputer	Q45.2B	2	115866.53	7033.77	0.9	26.2	0	0

To view more information about a host, click the name of the host. The Overview tab of the Hosts page appears. For more information, see [Host overview](#).

3.1.7.4.5 Cluster scale-out

This topic describes how to add physical hosts to scale out a DataHub cluster in Apsara Bigdata Manager (ABM). Cluster scale-out refers to the process of adding physical hosts in the default cluster of Apsara Infrastructure Management Framework to a DataHub cluster. The physical hosts of a DataHub cluster include chunkserver and frontend hosts.

Prerequisites

- The physical hosts to be added to a DataHub cluster are available in the default cluster of Apsara Infrastructure Management Framework.
- The default cluster of Apsara Infrastructure Management Framework has hosts whose product type is datahub.



Note:

Currently, scale-out is only available for chunkserver and frontend nodes in a DataHub cluster.

Context

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to the default cluster of Apsara Infrastructure Management Framework. The default

cluster can be considered as an idle resource pool that provides resources for scaling out clusters for your business. To scale out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

Step 1: Obtain the name of the host to be added to a DataHub cluster

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click TIANJI to log on to the Apsara Infrastructure Management Framework console.
3. In the top navigation bar of the page that appears, choose Operations > Machine Operations.
4. On the Machine Operations page that appears, search for a host whose product type is datahub in the default cluster. Copy the name of the host.

Step 2: Add the host to a DataHub cluster

You can add multiple hosts to a DataHub cluster at a time to scale out the cluster. To achieve this, you need to first specify an existing host as the template host. When you scale out the DataHub cluster, the hosts copy configurations from the template host so that the hosts can be added to the cluster at a time.

1. On the O&M page of the ABM console, click the Clusters tab. On the page that appears, select a cluster in the left-side navigation pane. Click the Hosts tab, and then select a host whose role is chunkserver or frontend as the template host.
2. Choose Actions > Scale out Cluster. In the Scale out Cluster dialog box that appears, set relevant parameters.

The parameters are described as follows:

- **Refer Hostname:** the name of the template host. The name of the selected host is used by default.
 - **Hostname:** the name of the host to be added to the DataHub cluster. The drop-down list displays all available hosts in the default cluster for scale-out. You can select one or more hosts from the drop-down list.
3. Click Run. A message appears, indicating that the action has been submitted.

4. View the scale-out status.

Click **Actions** in the upper-left corner, and then click **Execution History** next to **Scale out Cluster** to view the scale-out history.

It may take some time for the cluster to be scaled out. In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

5. If the status is **RUNNING**, click **Details** to view the steps and progress of the scale-out.

6. If the status is **FAILED**, click **Details** to locate the failure cause. For more information, see [Locate failure causes](#).

Locate failure causes

1. On the **Clusters** page, click **Actions**, and then click **Execution History** next to **Scale out Cluster** to view the scale-in history.

2. If the status of a record is **FAILED**, click **Details** to locate the failure cause.

You can also view the parameter settings, host details, script, and execution parameters to locate the failure cause.

3.1.7.4.6 Cluster scale-in

This topic describes how to remove physical hosts to scale in a DataHub cluster in Apsara Bigdata Manager (ABM). Cluster scale-in refers to the process of removing physical hosts from a DataHub cluster to the default cluster of Apsara Infrastructure Management Framework. The physical hosts of a DataHub cluster include chunkserver and frontend hosts.

Prerequisites

- Currently, scale-in is only available for chunkserver and frontend nodes in a DataHub cluster.
- Before you remove one or more chunkserver nodes:
 - Run the `df` command to check the disk usage on each host. Calculate whether the disk will be full after a certain number of hosts are removed. If so, we recommend that you do not perform the scale-in.
 - Shards on the removed hosts will be migrated to other hosts. Therefore, you need to log on to the WebConsole to calculate the shard load on each host after the scale-in. If the number of shards on a host exceeds 1,000, performance

may be affected. In this case, we recommend that you do not perform the scale-in.

- **Before you remove one or more frontend nodes:**
 - **Run the `df` command to check the disk usage on each host. Calculate whether the disk will be full after a certain number of hosts are removed. If so, we recommend that you do not perform the scale-in.**
 - **Shards on the removed hosts will be migrated to other hosts. Therefore, you need to log on to the WebConsole to calculate the shard load on each host after the scale-in. If the number of shards on a host exceeds 1,000, performance may be affected. In this case, we recommend that you do not perform the scale-in.**
 - **Check the traffic and queries per second (QPS). If the traffic exceeds 400 MBit/s or the QPS exceeds 15,000, we recommend that you do not perform the scale-in.**

Context

In Apsara Stack, scaling out a cluster involves complex operations. You need to configure a new physical host on Deployment Planner so that it can be added to the default cluster of Apsara Infrastructure Management Framework. The default cluster can be considered as an idle resource pool that provides resources for scaling out clusters for your business. To scale out a cluster, add physical hosts in the default cluster of Apsara Infrastructure Management Framework to the cluster. To scale in a cluster, remove physical hosts from the cluster to the default cluster of Apsara Infrastructure Management Framework.

Procedure

- 1. On the O&M page of the ABM console, click the Clusters tab. On the page that appears, select a cluster in the left-side navigation pane. Click the Hosts tab, and then select one or more hosts whose role is chunkserver or frontend.**
- 2. On the Clusters page, choose Actions > Scale in Cluster. The Scale in Cluster dialog box appears.**

Hostname: the name of the host to be removed from the DataHub cluster. The name of the selected host is used by default.

- 3. Click Run. A message appears, indicating that the action has been submitted.**

4. View the scale-in status.

Click **Actions** in the upper-left corner, and then click **Execution History** next to **Scale in Cluster** to view the scale-in history.

It may take some time for the cluster to be scaled in. In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.



Note:

If the status is **FAILED**, click **Details** to locate the failure cause. For more information, see [Locate failure causes](#).

5. View the scale-in progress.

If the status is **RUNNING**, click **Details** to view the steps and progress of the scale-in.

Locate failure causes

1. On the **Clusters** page, click **Actions**, and then click **Execution History** next to **Scale in Cluster** to view the scale-in history.
2. If the status of a record is **FAILED**, click **Details** to locate the failure cause.

You can also view the parameter settings, host details, script, and execution parameters to locate the failure cause.

3.1.7.4.7 Delete topics from a smoke testing project

Apsara Bigdata Manager (ABM) allows you to delete topics from a DataHub test project and view the execution history.

1. On the **Clusters** page, select a cluster in the left-side navigation pane. Click the **Hosts** tab. The **Hosts** page for the cluster appears.
2. On the **Clusters** page, choose **Actions > Delete Topic from Smoke Testing**. The **Delete Topic from Smoke Testing** dialog box appears.
3. Click **Run**. A message appears, indicating that the action has been submitted.

4. View the history of deleting topics.

Click **Actions** in the upper-left corner, and then click **Execution History** next to **Delete Topic from Smoke Testing** to view the execution history.

In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

- If the status is **FAILED**, click **Details** to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

- If the status is **SUCCESS**, click **Details** to view the execution result. On the page that appears, click **View Details** in the **Actions** column. The execution result including the time when the job was run and the IP address of the host appears in the **Execution Details** section in the lower-right corner.

3.1.7.4.8 Reverse parse RequestId

Apsara Bigdata Manager (ABM) allows you to reverse parse RequestId in DataHub to obtain the time when a job was run and the IP address of the host. You can use the obtained information to query logs for troubleshooting.

1. On the **Clusters** page, select a cluster in the left-side navigation pane. Click the **Hosts** tab. The **Hosts** page for the cluster appears.
2. On the **Clusters** page, choose **Actions > Reverse Parse Request ID**. In the **Reverse Parse Request ID** dialog box that appears, set **Request Id**.
3. Click **Run**. A message appears, indicating that the action has been submitted.

4. View the reverse parsing status.

Click **Actions** in the upper-left corner, and then click **Execution History** next to **Reverse Parse Request ID** to view the execution history.

In the **Current Status** column, **RUNNING** indicates that the execution is in progress, **SUCCESS** indicates that the execution is successful, and **FAILED** indicates that the execution fails.

- If the status is **FAILED**, click **Details** to locate the failure cause.

You can also view information about parameter settings, host details, script, and execution parameters to locate the failure cause.

- If the status is **SUCCESS**, click **Details** to view the execution result. On the page that appears, click **View Details** in the **Actions** column. The execution result including the time when the job was run and the IP address of the host appears in the **Execution Details** section in the lower-right corner.

3.1.7.5 Host O&M

3.1.7.5.1 Host O&M overview

This topic describes the features of DataHub host O&M and how to access the host O&M page.

Features

- **Overview page:** displays information about hosts in a DataHub cluster, including the host information, service role status, health check status, health check history, and trend charts of the CPU, memory, storage, load, and packet loss rate metrics.
- **Charts page:** displays the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.
- **Health Status page:** displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.
- **Services page:** displays the cluster, service instances, and service instance roles of a host.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click **DataHub**.

3. On the DataHub page that appears, click O&M in the upper-right corner, and then click the Hosts tab.
4. On the Hosts page, select a host in the left-side navigation pane. The Overview for the host appears.

3.1.7.5.2 Host overview

The host overview page displays the overall running information about a host in a DataHub cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

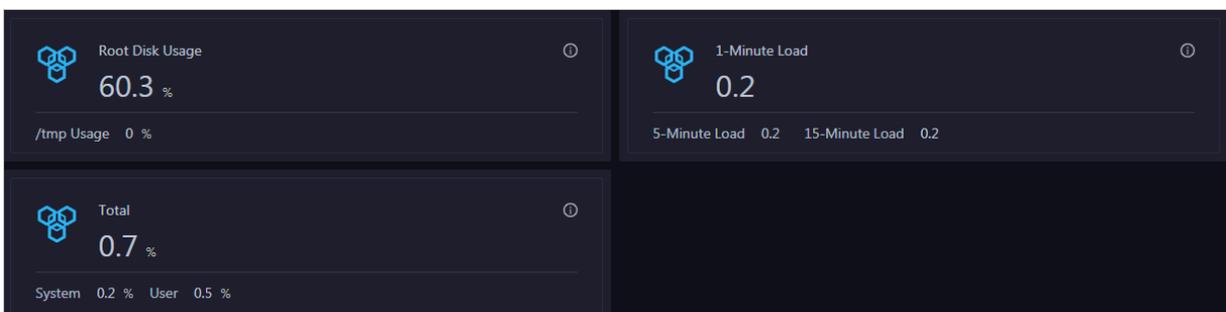
Entry

On the Hosts page, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.

On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Root Disk Usage, Total, and 1-Minute Load

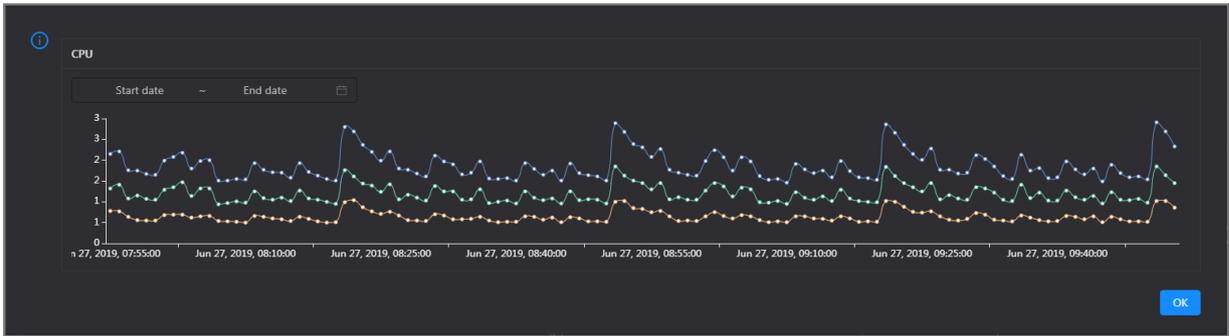
These sections display the root disk usage, total usage, and 1-minute load for the selected host. The Root Disk Usage section provides the usage of the `/tmp` directory. The Total section provides the system usage and user usage. The 1-Minute Load section provides the 1-minute, 5-minute, and 15-minute load averages.



CPU

This chart displays the trend lines of the total CPU usage (`cpu`), CPU usage for executing code in kernel space (`sys`), and CPU usage for executing code in user space (`user`) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

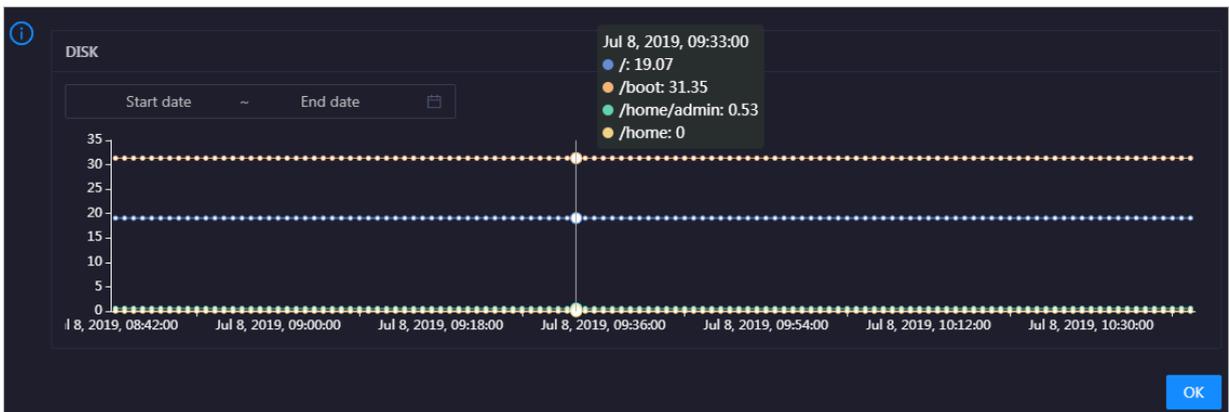


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

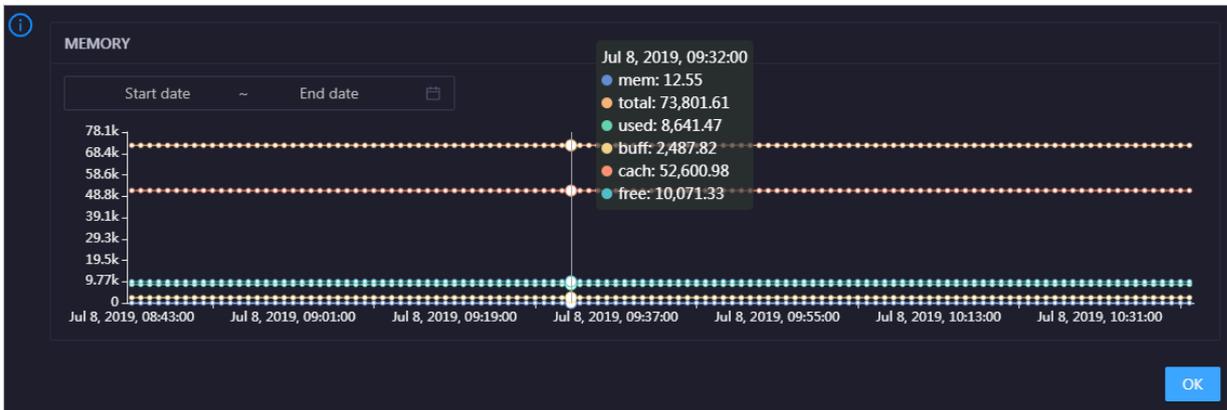


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (`mem`), total memory size (`total`), used memory size (`used`), size of memory used by kernel buffers (`buff`), size of memory used by the page cache (`cach`), and available memory size (`free`) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

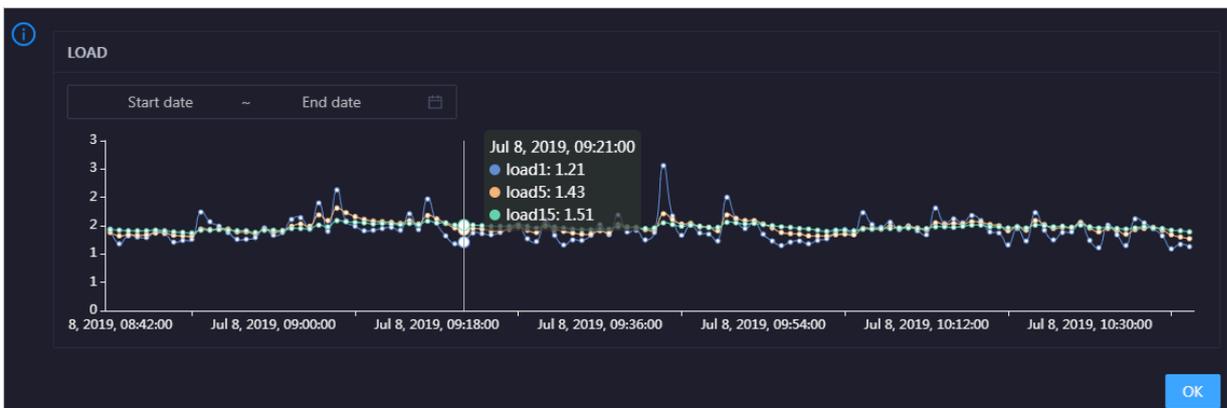


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

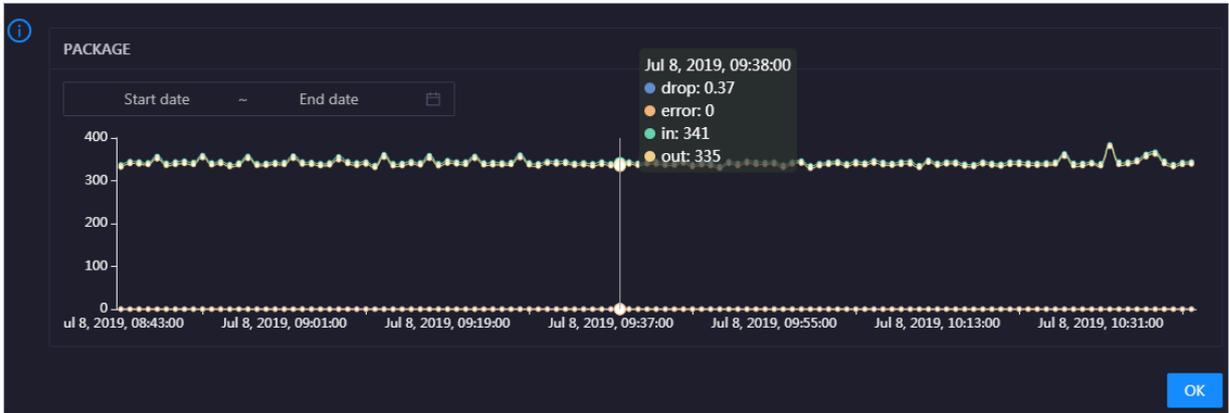


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.

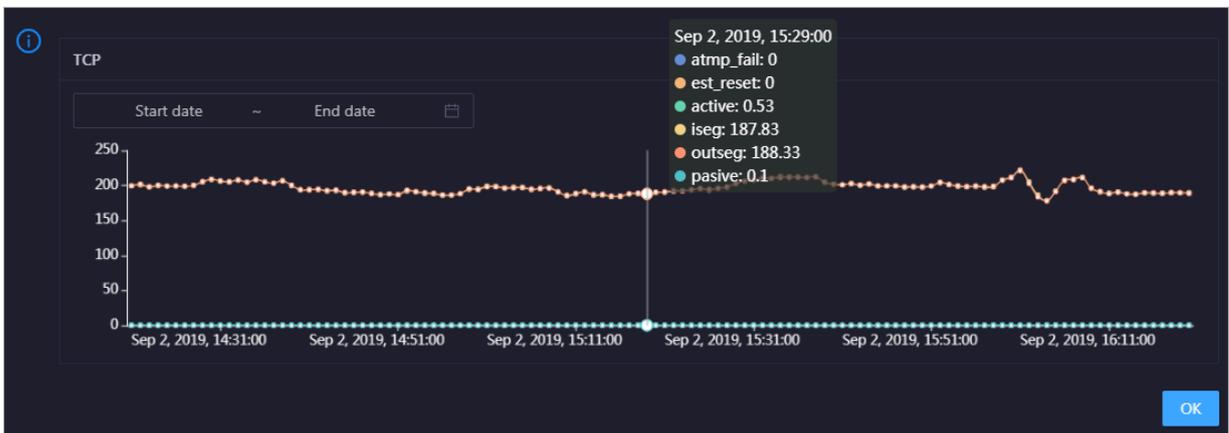


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click  in the upper-right corner of the chart to zoom in the chart.

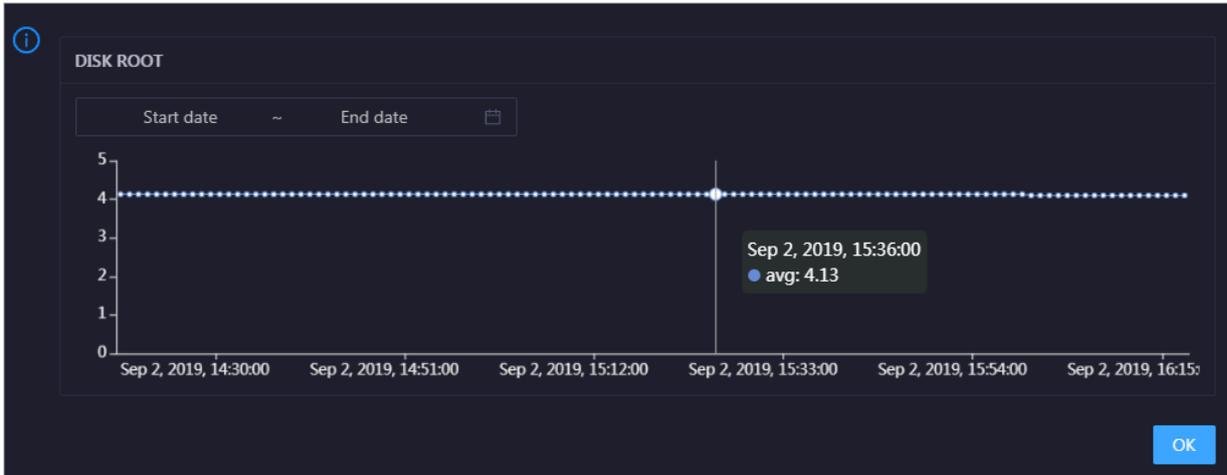


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the host over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

Click [View Details](#) to go to the [Host health](#) page. On this page, you can view the health check details.

Health Check History

This section displays a record of the health checks performed on the host.

Health Check History		View Details
Time	Event Content	
Recently	1 alerts are reported by checkers.	< 1 >

Click [View Details](#) to go to the [Host health](#) page. On this page, you can view the health check details.

You can click the event content of a check to view the exception items.

Details				✕
Checker	Host	Status	Status Updated At	
bcc_host_live_check		CRITICAL	Jul 7, 2019, 18:35:30	< 1 >

3.1.7.5.3 Host charts

On the host chart page, you can view the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.

On the Hosts page, select a host in the left-side navigation pane, and then click the Charts tab. The Charts page for the host appears.



The Charts page displays trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host. For more information, see [Host overview](#).

3.1.7.5.4 Host health

On the host health status page, you can view the checkers of all hosts, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

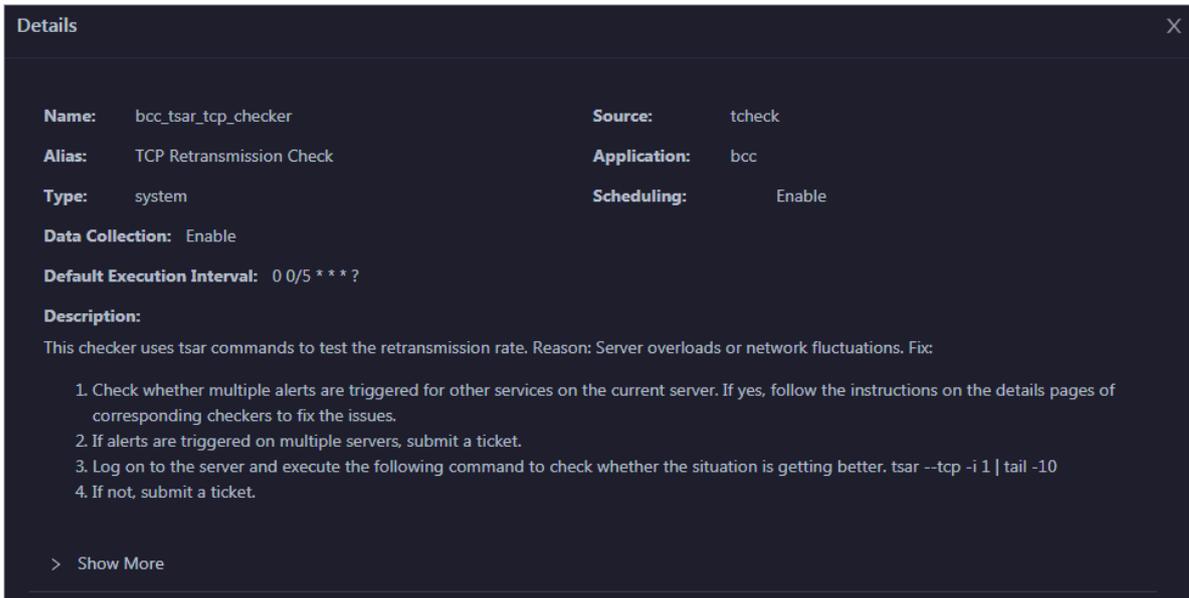
Entry

On the Hosts page, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.

On the Health Status page, you can view all checkers of the host and the check results for the hosts in the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

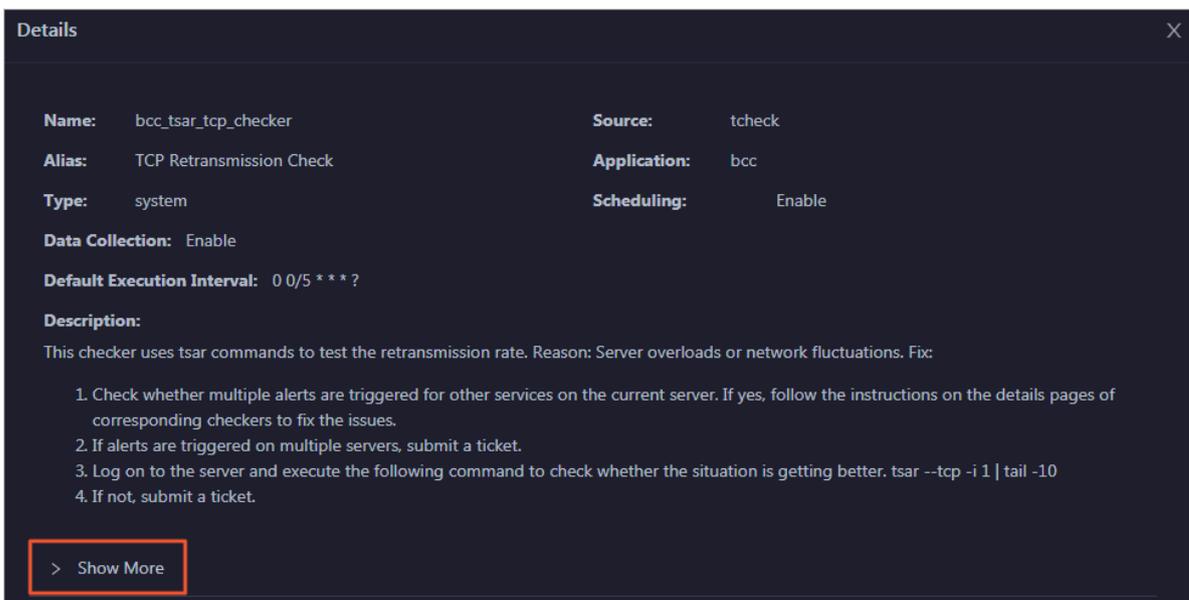
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

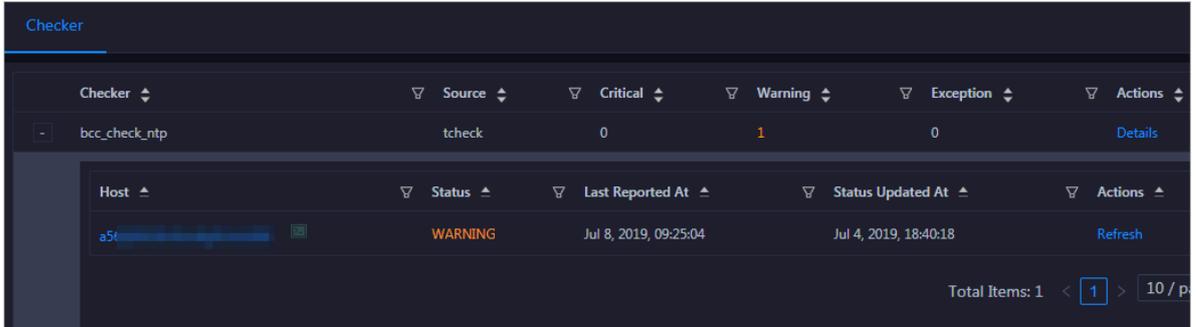


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

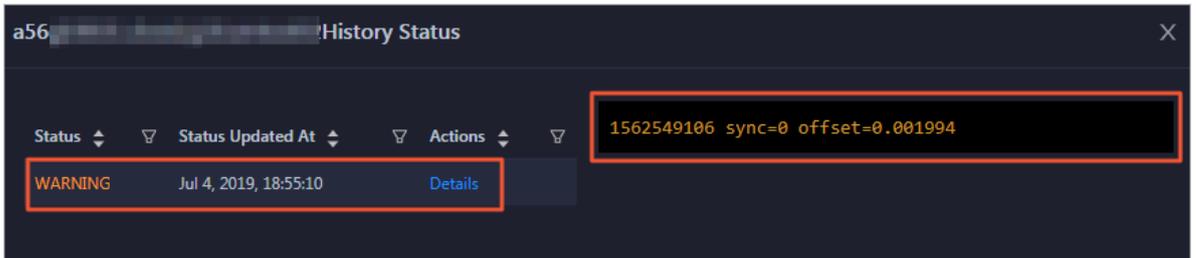
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

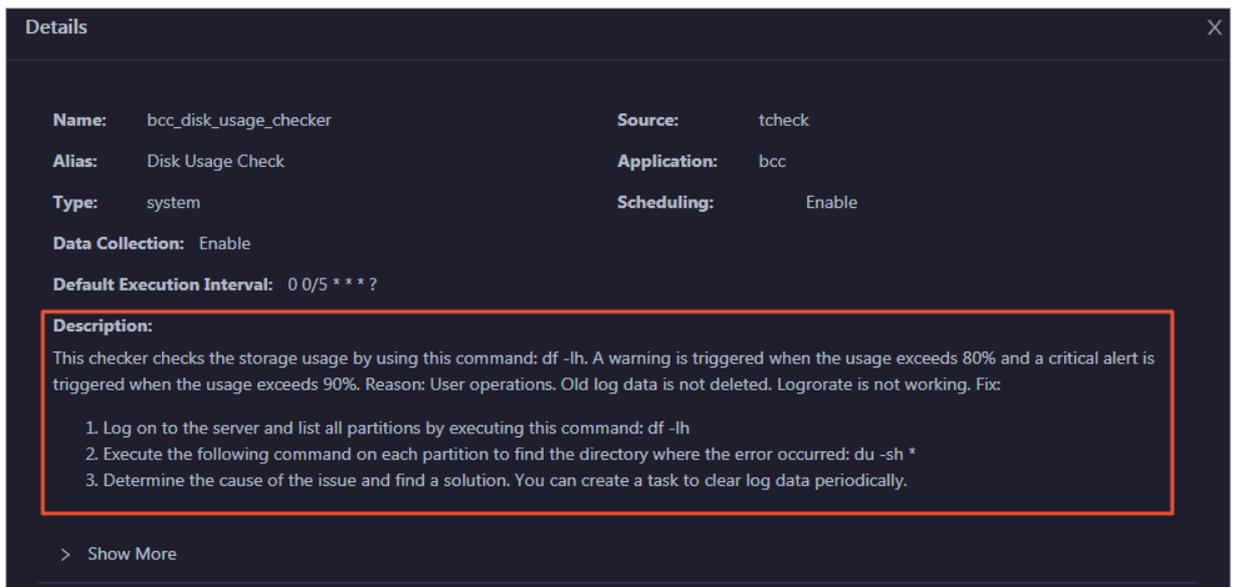


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

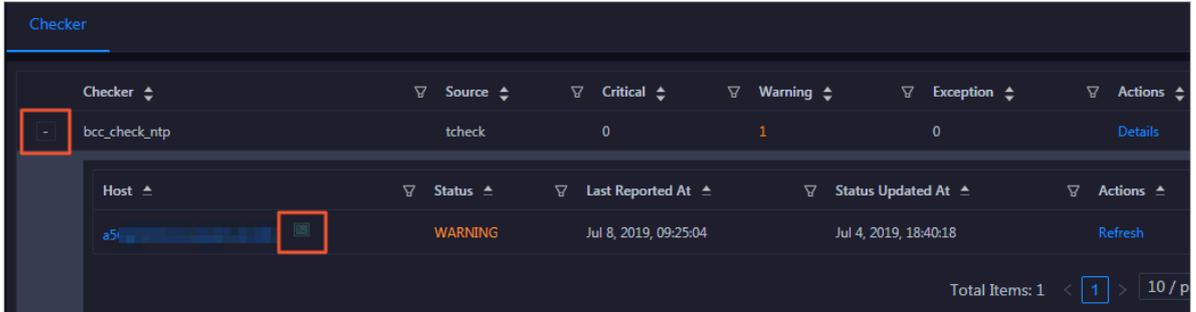
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



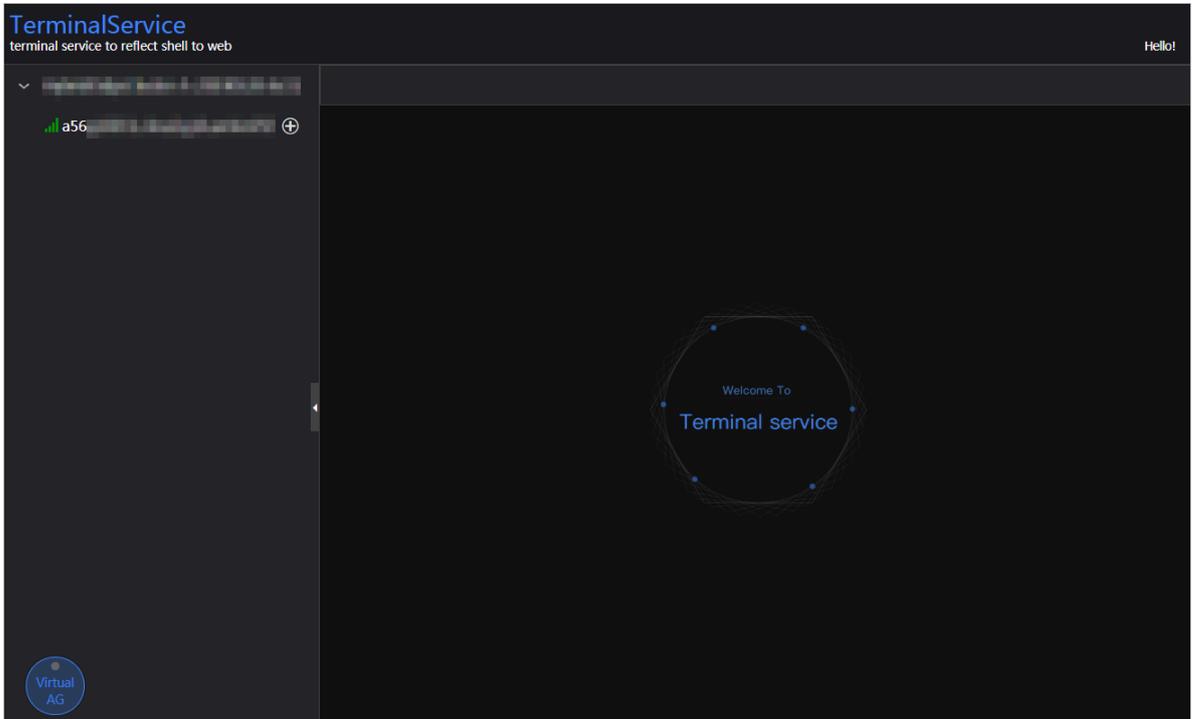
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

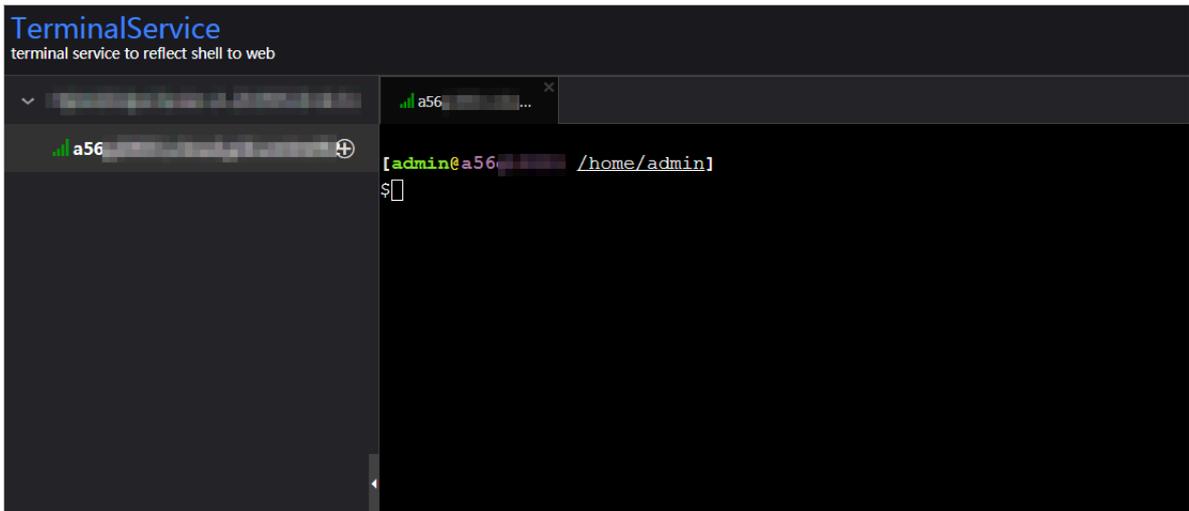
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

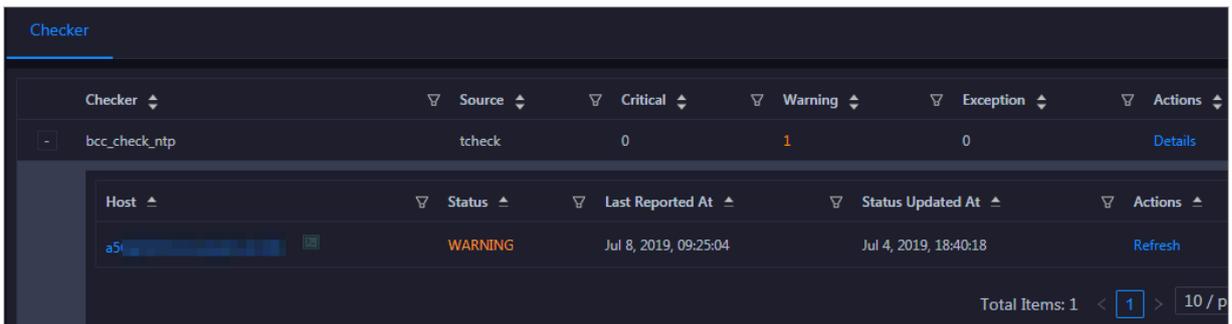


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.7.5.5 Host services

On the Services page, you can view information about service instances and service instance roles of a host.

On the Hosts page, select a host in the left-side navigation pane, and then click the Services tab. The Services page for the host appears.

On the Services page, you can view the cluster, service instances, and service instance roles of the host.

3.1.8 Elasticsearch

3.1.8.1 O&M overview

This topic describes the features of Elasticsearch O&M and how to access the Elasticsearch O&M page.

Modules

Elasticsearch O&M includes business O&M, service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Business O&M	Cluster Configuration	Allows you to view and modify the cluster configuration files of the worker and kibana nodes for Elasticsearch.
	System Configuration	Allows you to view and modify the system configuration files for Elasticsearch.
Service O&M	Overview	Displays all Elasticsearch services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.
	Hosts	Displays all hosts where each Elasticsearch service is run so that you can understand the service deployment on hosts.
Cluster O&M	Overview	Displays the overall running and health check information about a cluster. On this page, you can view the host status, service status, health check result, and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, heap memory usage, TCP connection, and root disk usage.
	Health Status	Displays all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Module	Feature	Description
Host O&M	Overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Charts	Displays the enlarged trend charts of CPU usage , memory usage, storage usage, load, and packet transmission.
	Health Status	Displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.
	Services	Displays information about service instances and service instance roles of a host.

Entry

1. Log on to the ABM console.
2. Click  in the upper-left corner, and then click Elasticsearch.
3. On the Elasticsearch page that appears, click O&M in the upper-right corner. The Business page appears.

The O&M page includes four modules, namely, Business, Services, Clusters, and Hosts.

3.1.8.2 Business O&M

3.1.8.2.1 Cluster configuration

This topic describes how to view and modify the cluster configuration files of the worker and kibana nodes for Elasticsearch in Apsara Bigdata Manager (ABM).

Entry

1. At the top of the O&M page, click the Business tab.
2. On the Business page that appears, click Cluster Configuration in the left-side navigation pane.

3. In the worker or kibana list, click a cluster configuration file that you want to view. The details of the file appear on the right.

Modify a cluster configuration file

1. Click a cluster configuration file to be modified and click Edit to modify the configuration file.
2. Click Save.
3. Click Preview.
 - a. In the Preview dialog box that appears, you can compare the differences before and after the file modification.
 - b. If the modification is correct, click OK.
4. Click Submit at the bottom of the page. The modification is completed.

If you want to undo the modification, click Undo.

Upload a plug-in



Notice:

The custom plug-in may affect the stability of the cluster. Make sure that the custom plug-in is reliable and secure to use. The plug-in is not automatically updated with Elasticsearch. To update the plug-in, you must manually upload a new version of the plug-in.

1. Select a cluster to which you want to upload a plug-in from the drop-down list. Click Upload Plug-in.
2. In the Upload Plug-in dialog box that appears, click Click here to select files for upload to upload one or more files.

To delete a file that no longer needs to be uploaded, click x next to the file.

3.



Notice:

The custom plug-in may affect the stability of the cluster. Make sure that the custom plug-in is reliable and secure to use. The plug-in is not automatically updated with Elasticsearch. To update the plug-in, you must manually upload a new version of the plug-in.

Select the check box in the dialog box.

4. Click OK.

3.1.8.2.2 System configuration

This topic describes how to view and modify the system configuration files for Elasticsearch in Apsara Bigdata Manager (ABM).

Entry

1. At the top of the O&M page, click the **Business** tab.
2. On the **Business** page that appears, click **System Configuration** in the left-side navigation pane.
3. Click a configuration file that you want to view. The details of the file appear on the right.

Modify a system configuration file

1. Click a system configuration file to be modified and click **Edit** to modify the configuration file.
2. Click **Save**.
3. Click **Preview**.
 - a. In the **Preview** dialog box that appears, you can compare the differences before and after the file modification.
 - b. If the modification is correct, click **OK**.
4. Click **Submit** at the bottom of the page. The modification is completed.

If you want to undo the modification, click **Undo**.

3.1.8.3 Service O&M

3.1.8.3.1 Service overview

The service overview page lists all Elasticsearch services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.

Entry

1. At the top of the O&M page, click the **Services** tab.
2. On the **Services** page that appears, select a service in the left-side navigation pane. Click the **Overview** tab.

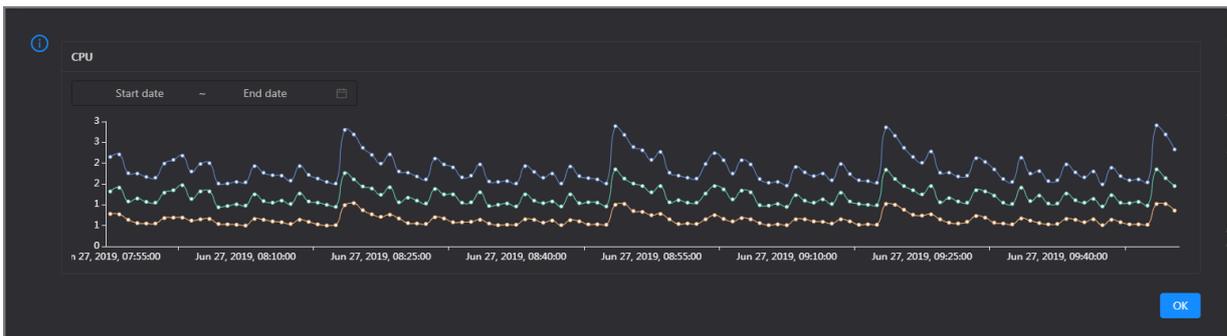
On the **Overview** page that appears, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

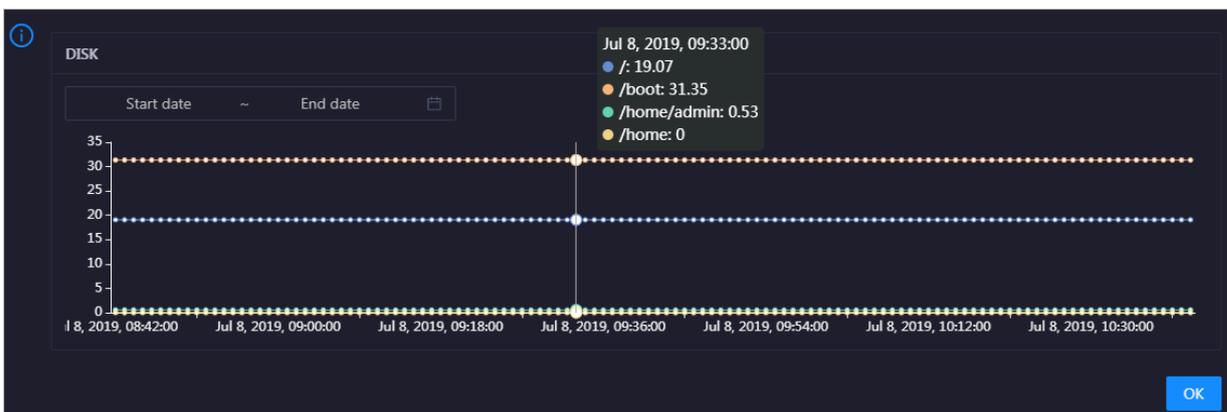
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the service in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

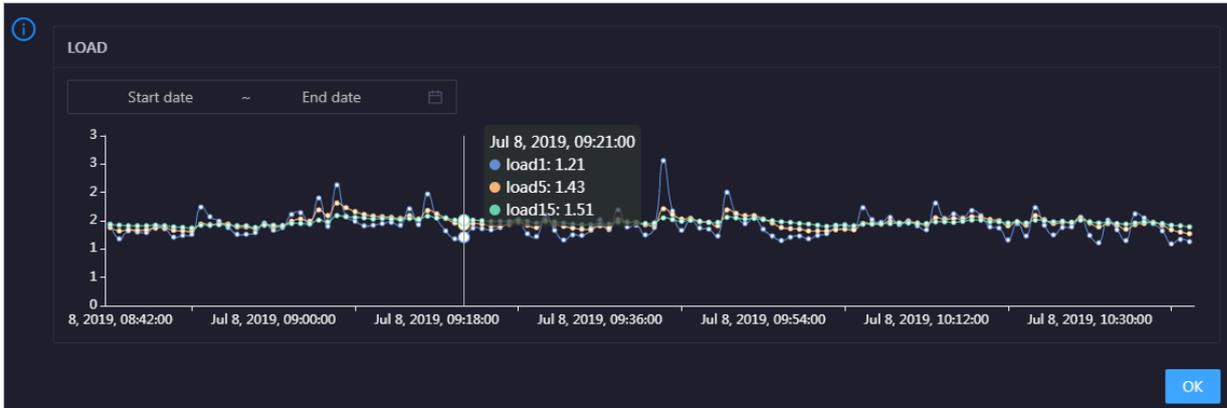


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

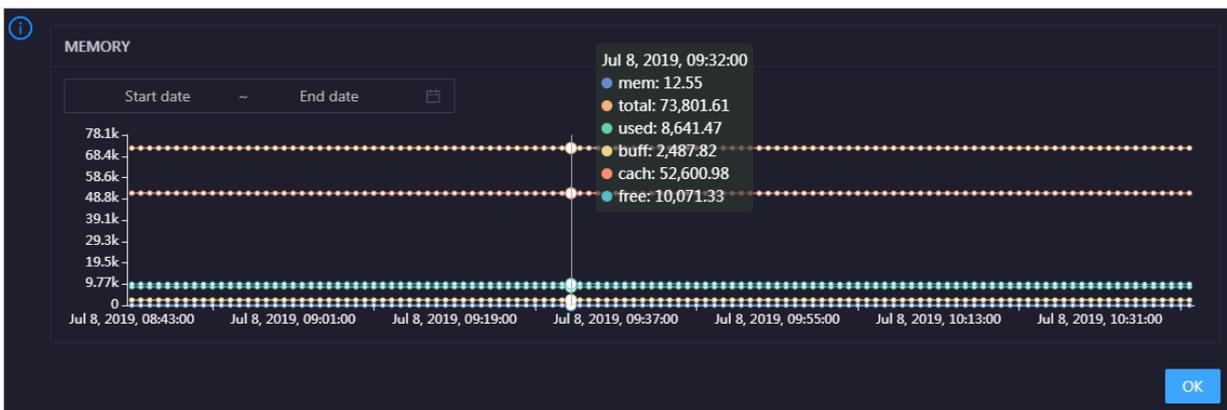


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



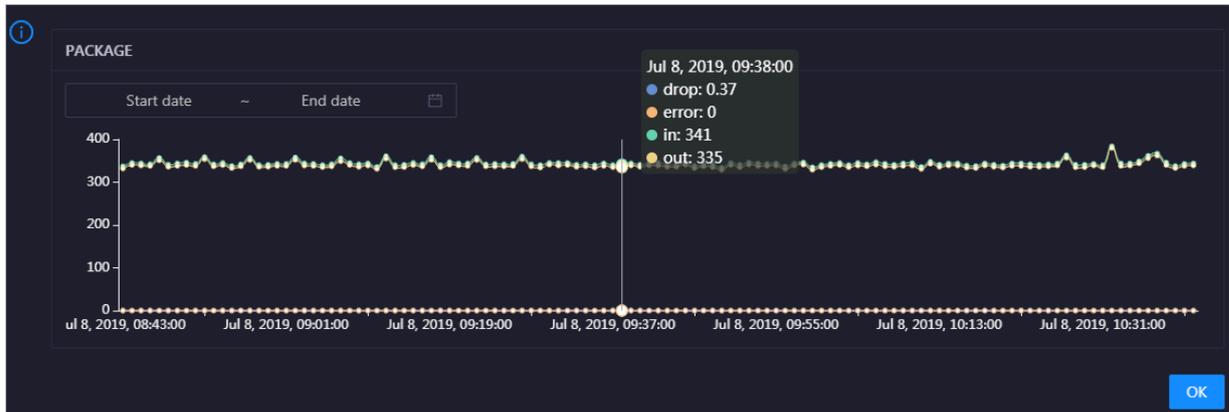
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for

the selected service over time in different colors. These trend lines reflect the data transmission status of the service.

Click  in the upper-right corner of the chart to zoom in it.

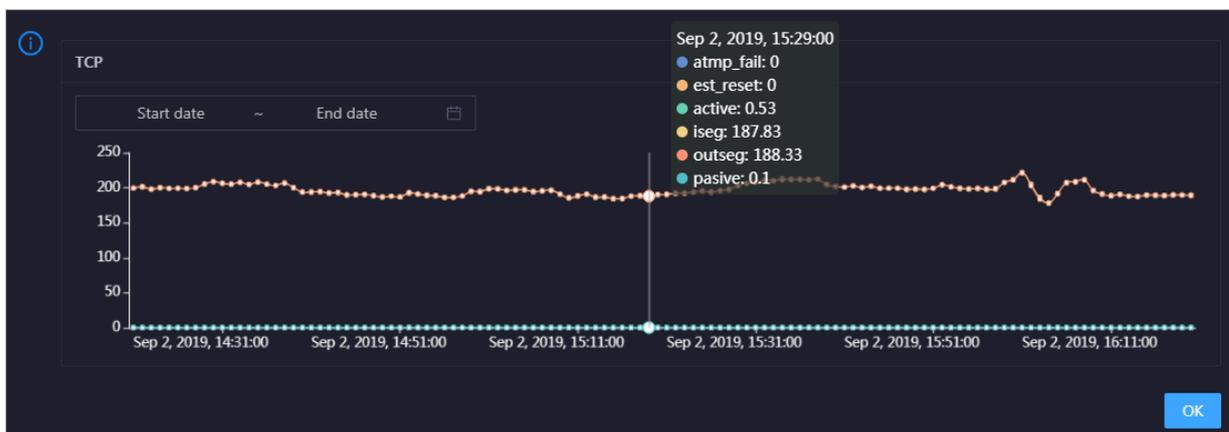


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the selected service over time in different colors. These trend lines reflect the TCP connection status of the service.

Click  in the upper-right corner of the chart to zoom in it.

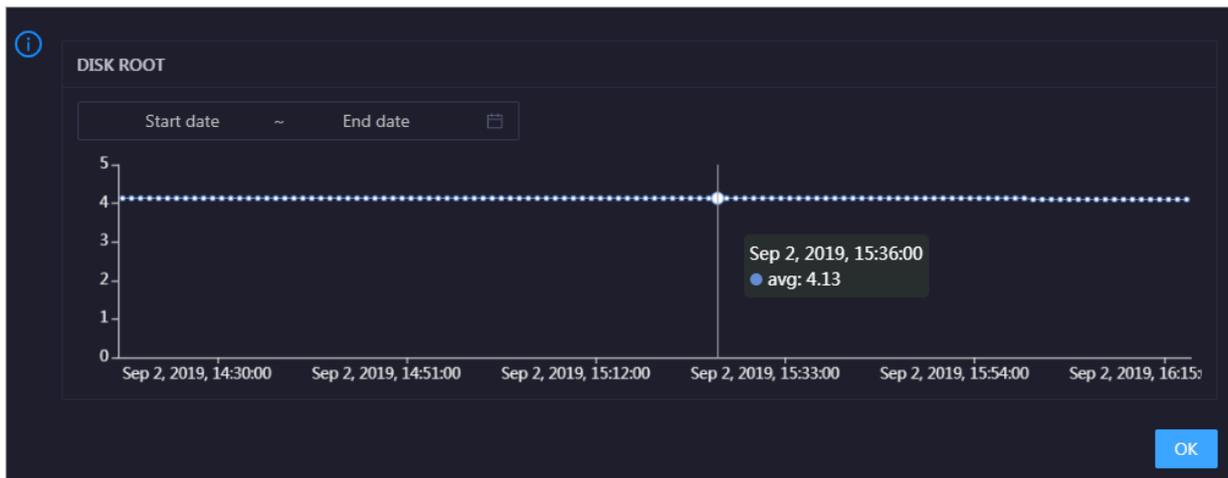


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

This chart displays the trend line of the average root disk usage (avg) for the selected service over time.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

3.1.8.3.2 Service hosts

This topic describes how to view all hosts where each Elasticsearch service is run.

On the Server page, you can view the hosts where the selected service is run.

1. At the top of the O&M page, click the Services tab.
2. On the Services page that appears, select a service in the left-side navigation pane.
3. Click the Server tab. The Server page for the service appears.

On the Server page, you can view the hosts where the selected service is run.

3.1.8.4 Cluster O&M

3.1.8.4.1 Cluster overview

The cluster overview page displays the overall running and health check information about a cluster. On this page, you can view the host status, service

status, health check result, and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, heap memory usage, TCP connection, and root disk usage.

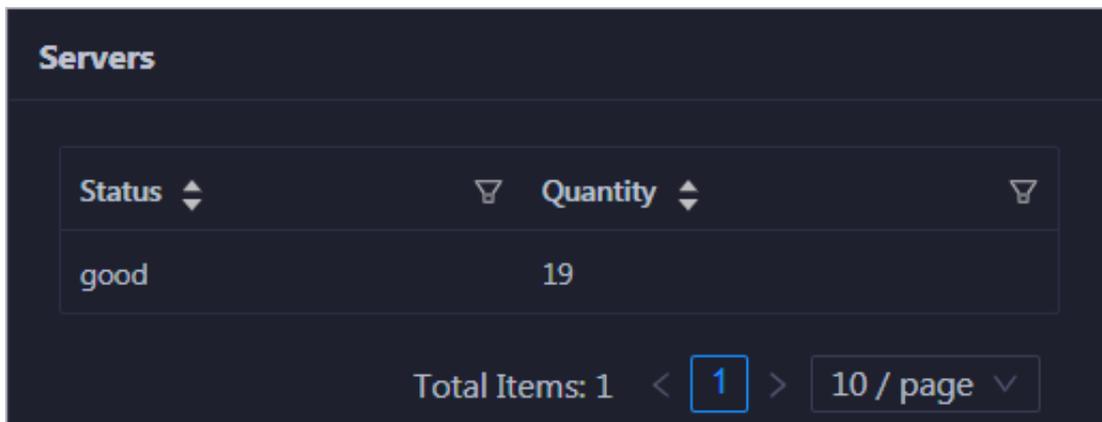
Entry

1. At the top of the O&M page, click the Clusters tab.
2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.

On the Overview page, you can view the host status, service status, health check result, and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, heap memory usage, TCP connection, and root disk usage. To view information about a cluster, select a region in the left-side navigation pane, and then select the cluster in the region.

Servers

This section displays all host statuses and the number of hosts in each status. The host statuses include good and bad.



Status	Quantity
good	19

Total Items: 1 < 1 > 10 / page

Services

This section displays all services deployed in the cluster and the respective number of available and unavailable services.

Health Check

This section displays the number of checkers deployed for the cluster and the respective number of Critical, Warning, and Exception alerts.

Health Check
View Details

Currently, 9 checkers are deployed on the service. 5 critical, 0 exception, and 0 warning alerts are reported.

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

Health Check History

This section displays a record of the health checks performed on the cluster.

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

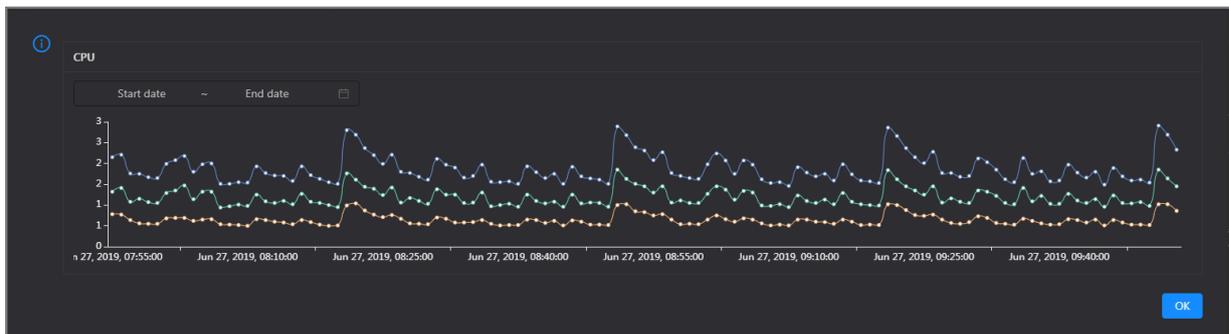
You can click the event content of a check to view the exception items.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

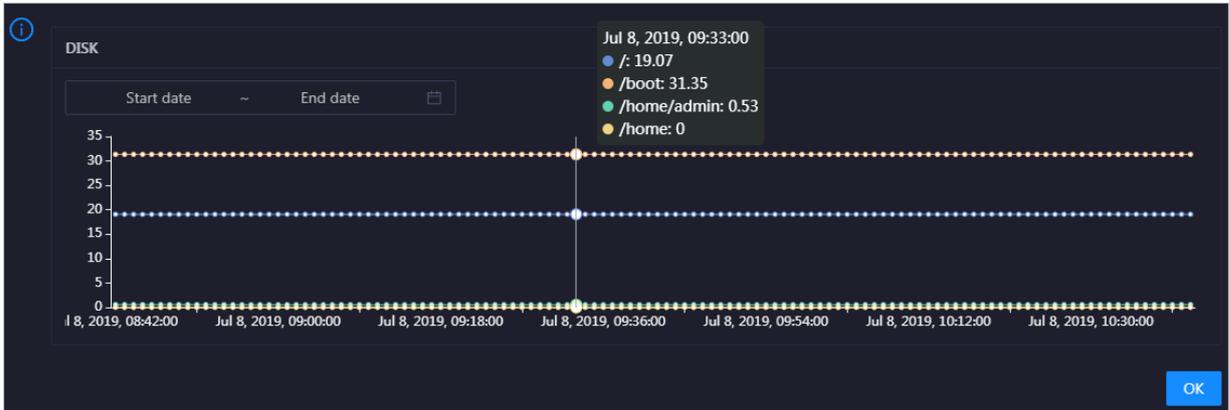
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

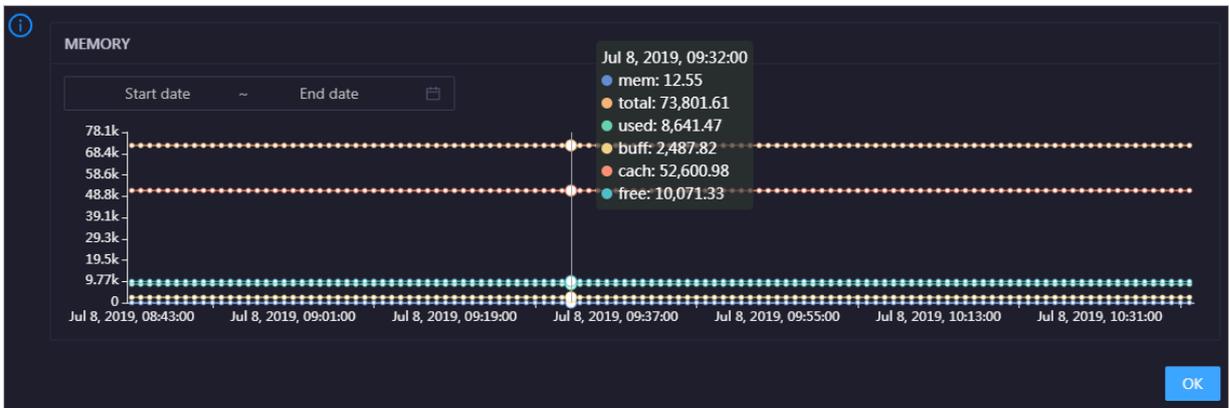


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

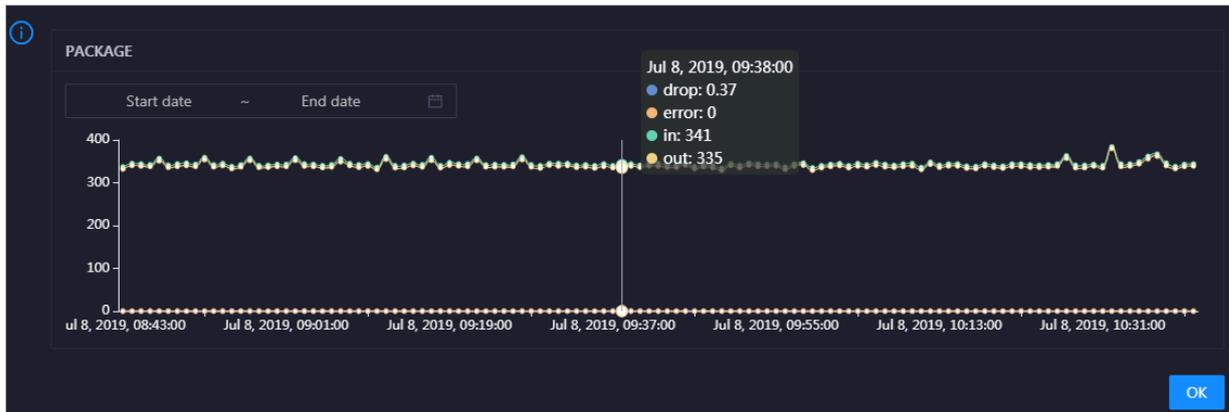


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

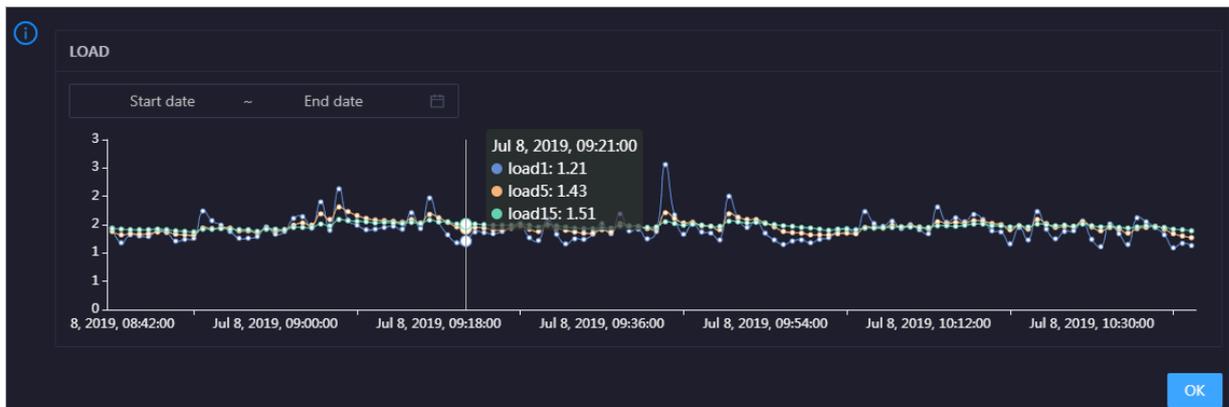


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

HEAP

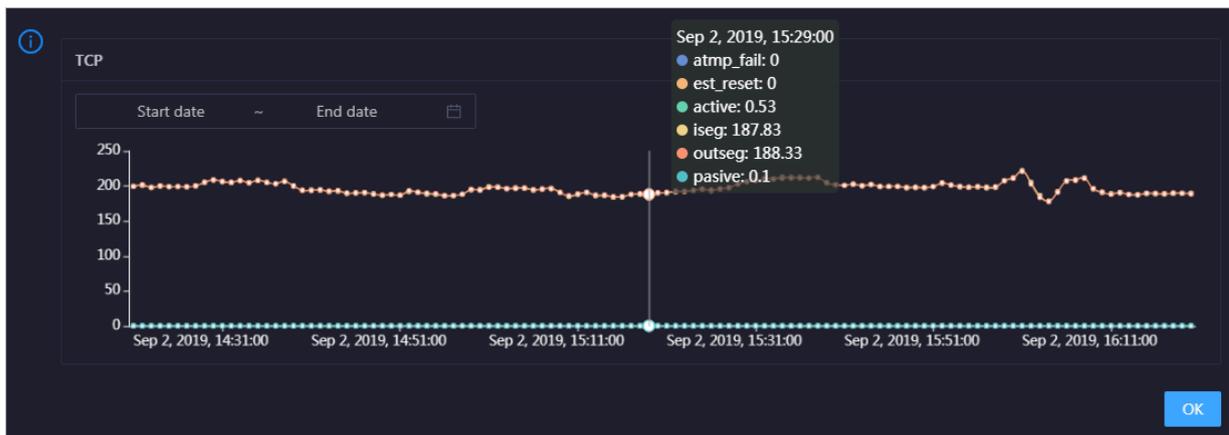
This chart displays the trend lines of the heap memory usage over time in different colors.

Click  in the upper-right corner of the chart to zoom in the chart.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

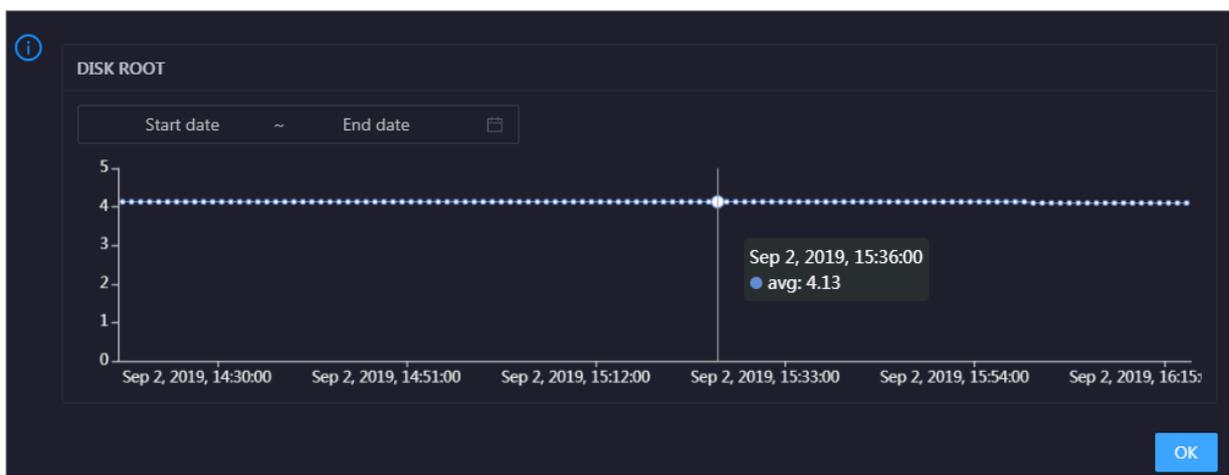


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

3.1.8.4.2 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

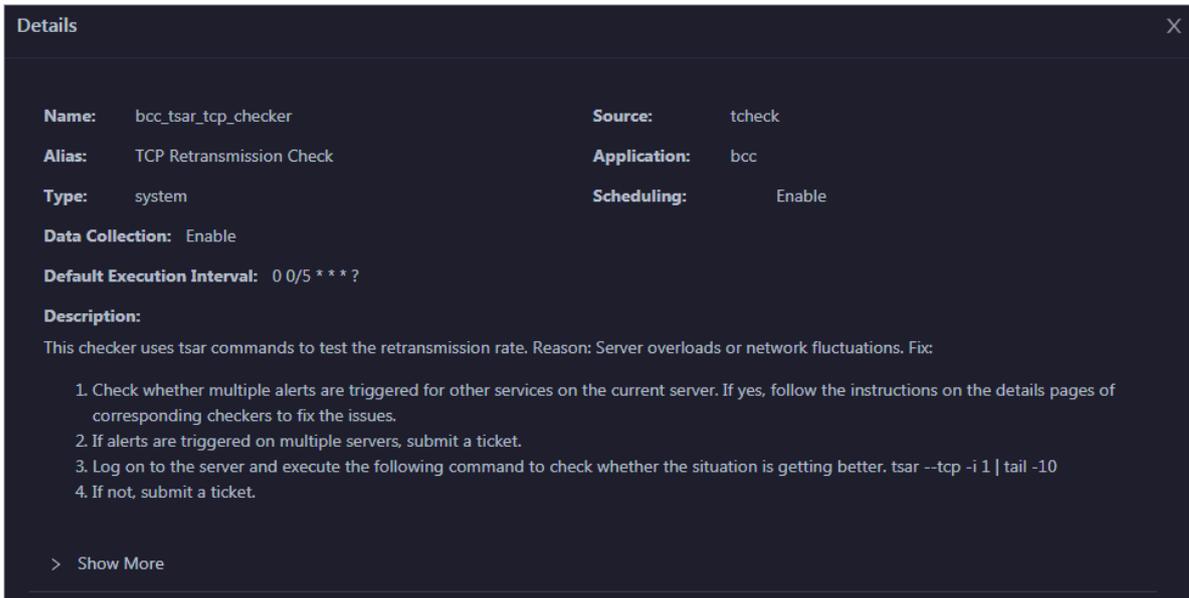
At the top of the O&M page, click the Clusters tab. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.

Checker	Source	Critical	Warning	Exception	Actions
+ bcc_host_live_check	tcheck	3	0	0	Details
+ elasticsearch_check_health_shuttle	tcheck	2	0	0	Details
+ bcc_check_ntp	tcheck	0	0	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_network_tcp_connections_checker	tcheck	0	0	0	Details
+ bcc_disk_usage_checker	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

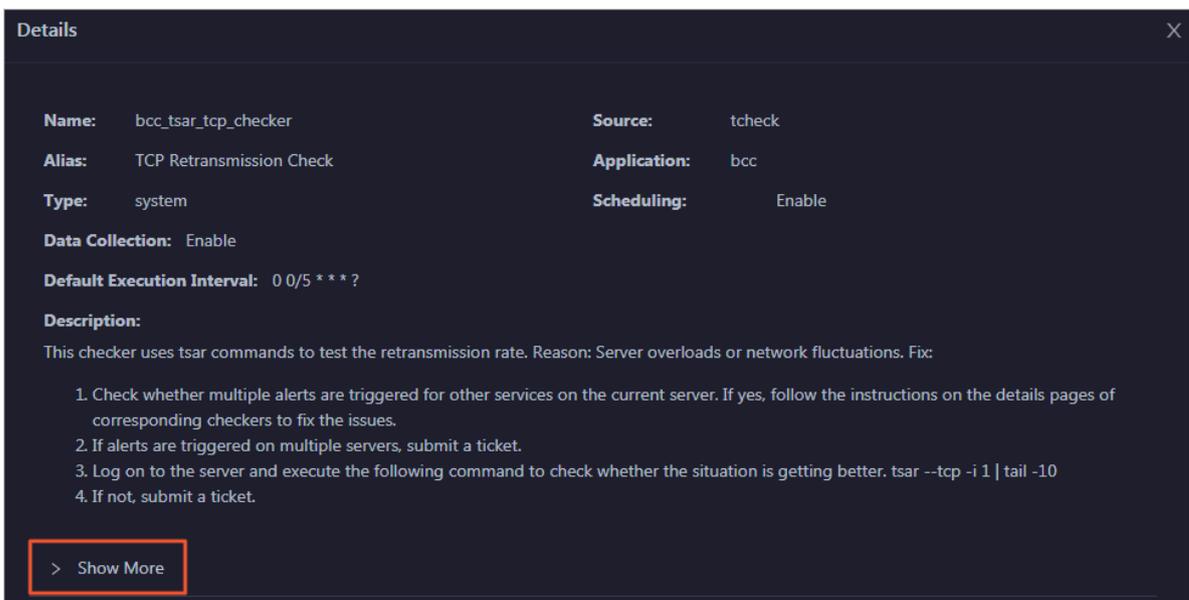
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

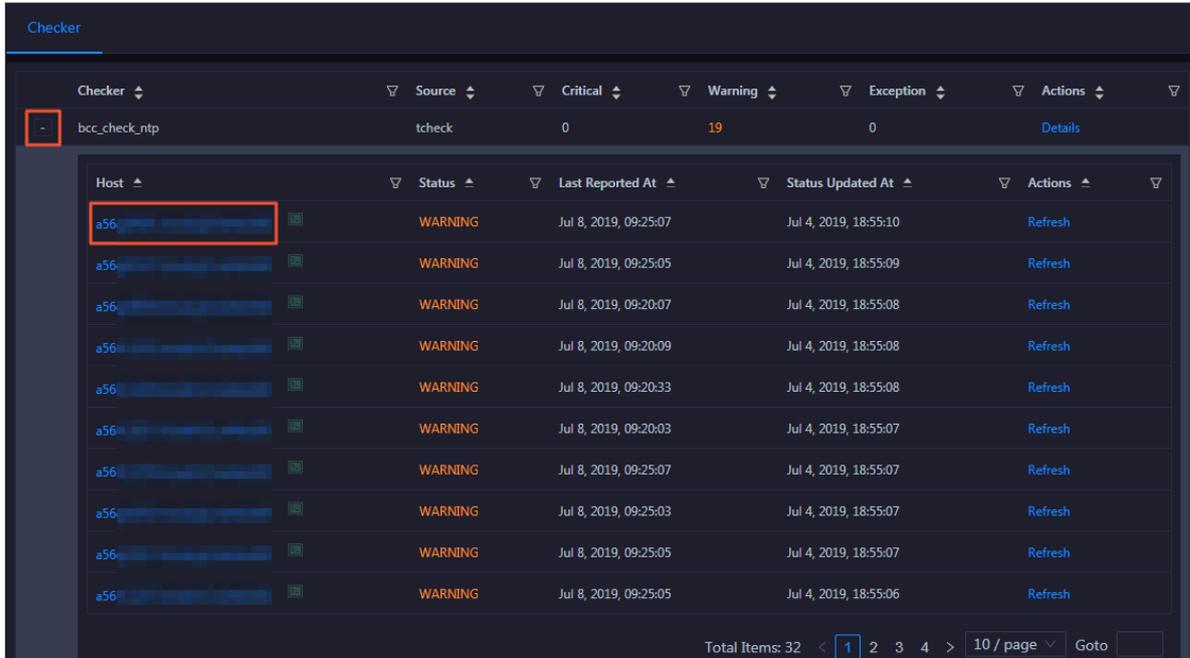


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

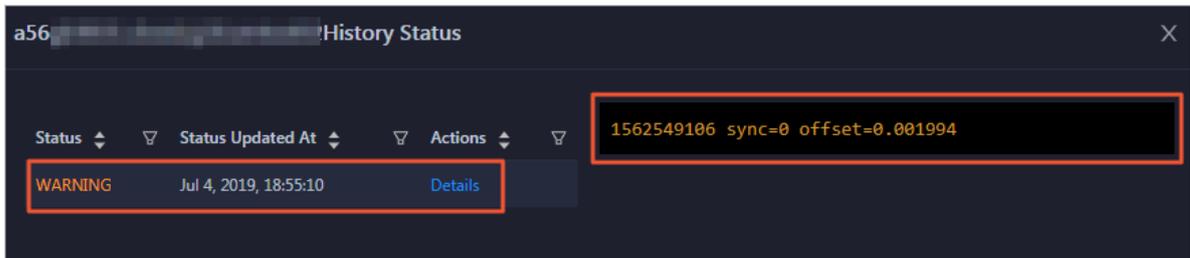
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

- 1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.**



- 2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**



Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

Details ✕

Name: bcc_disk_usage_checker	Source: tcheck
Alias: Disk Usage Check	Application: bcc
Type: system	Scheduling: Enable
Data Collection: Enable	
Default Execution Interval: 0 0/5 * * * ?	

Description:
 This checker checks the storage usage by using this command: df -lh. A warning is triggered when the usage exceeds 80% and a critical alert is triggered when the usage exceeds 90%. Reason: User operations. Old log data is not deleted. Logrotate is not working. Fix:

1. Log on to the server and list all partitions by executing this command: df -lh
2. Execute the following command on each partition to find the directory where the error occurred: du -sh *
3. Determine the cause of the issue and find a solution. You can create a task to clear log data periodically.

> Show More

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

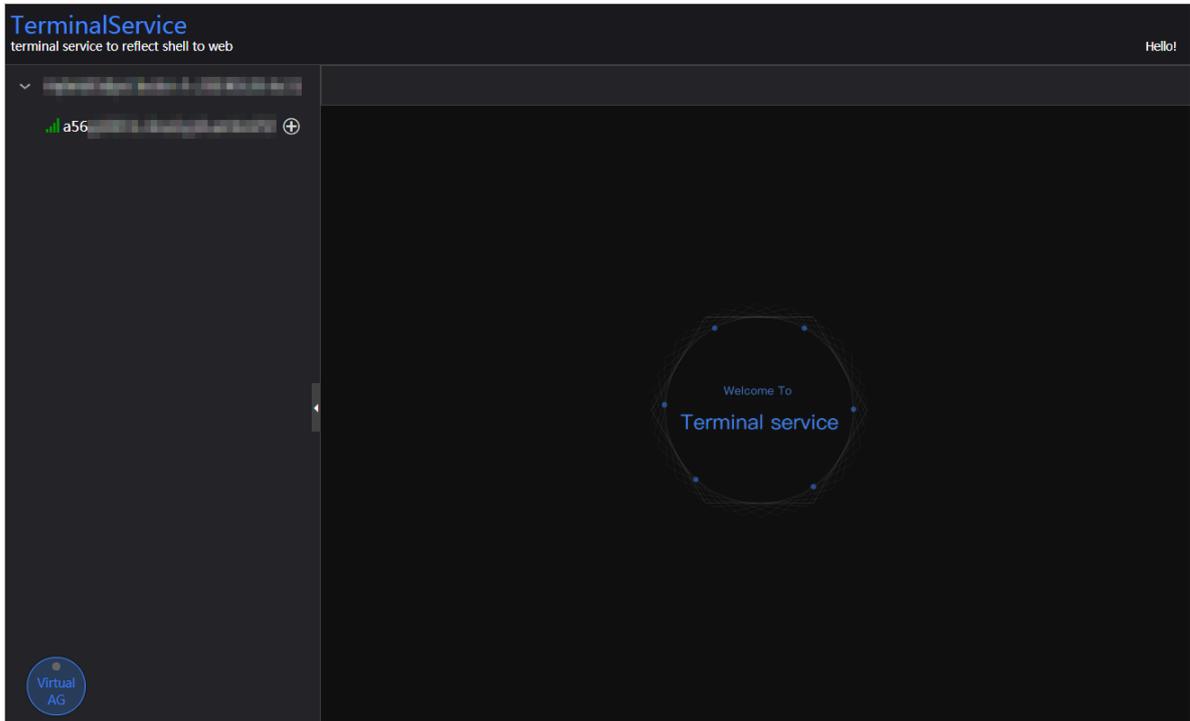
- 1. On the Health Status page, click + to expand a checker with alerts.**

Checker

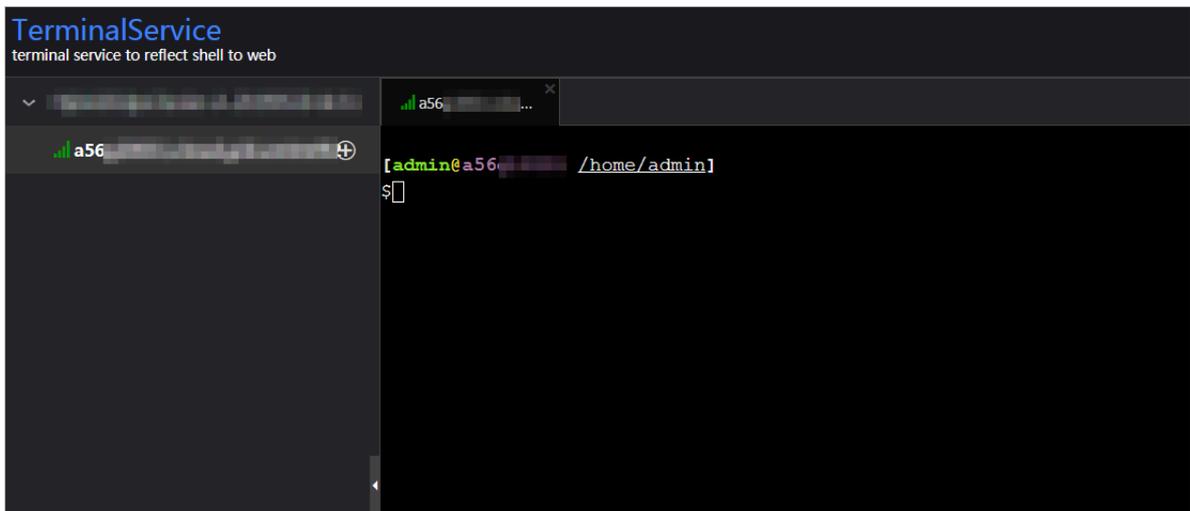
Checker	Source	Critical	Warning	Exception	Actions
- bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56 +	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56 +	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56 +	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.

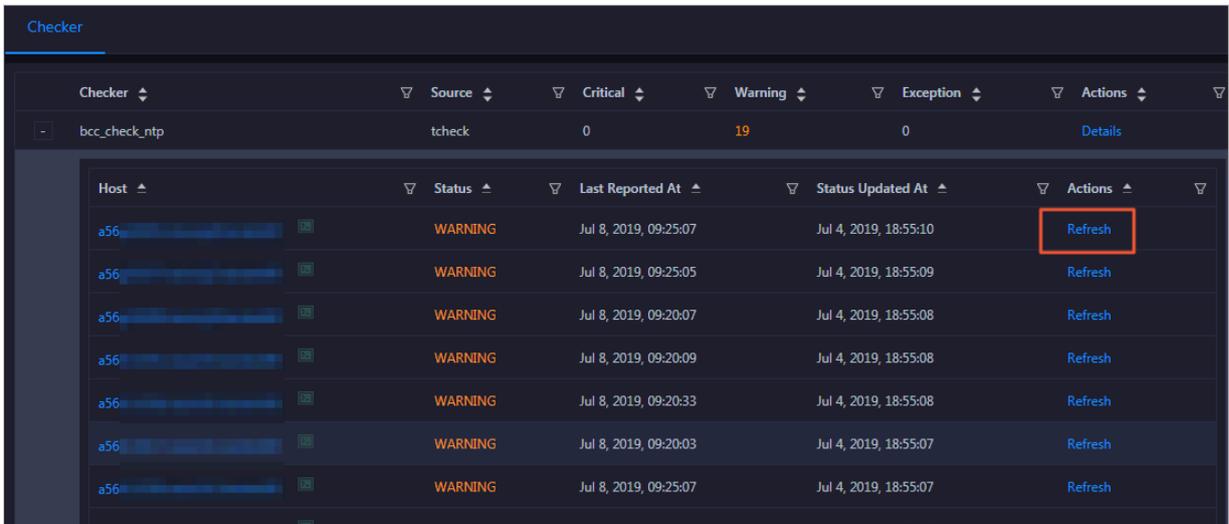


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



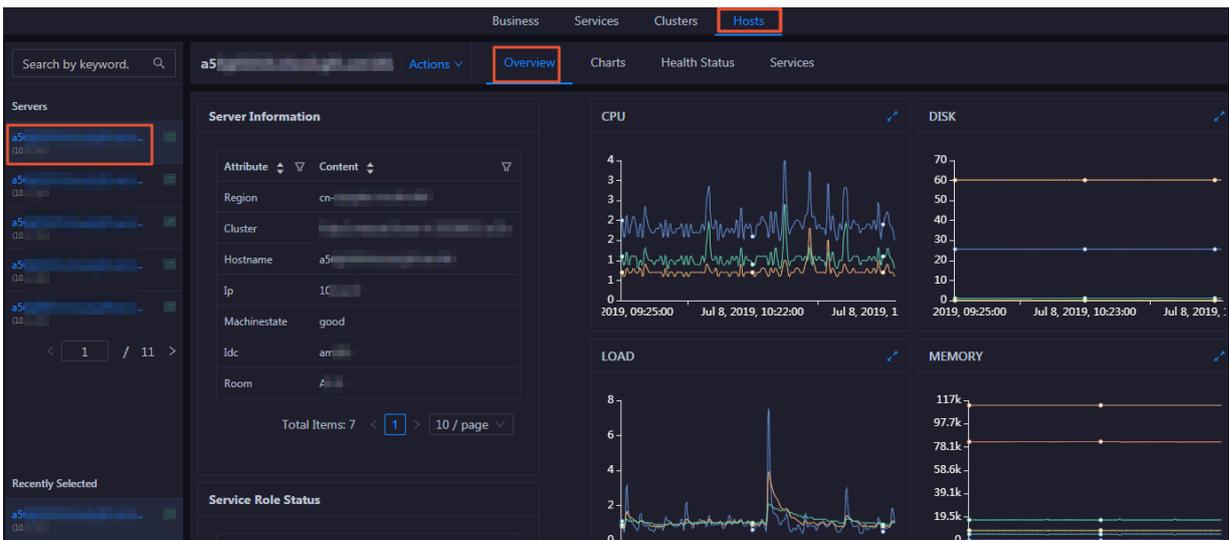
3.1.8.5 Host O&M

3.1.8.5.1 Host overview

The host overview page displays the overall running information about a host in an Elasticsearch cluster. On this page, you can view the information, service role status, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

Entry

On the Hosts page, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.



On the Overview page, you can view the server information, service roles, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

Server Information

This section displays the information about the host, including the region, cluster, name, IP address, status, IDC, and server room of the host.

Attribute	Content
Region	cn-...
Cluster	...
Hostname	a56...
Ip	10...
Machinestate	good
Idc	am...
Room	A...

Total Items: 7 < 1 > 10 / page

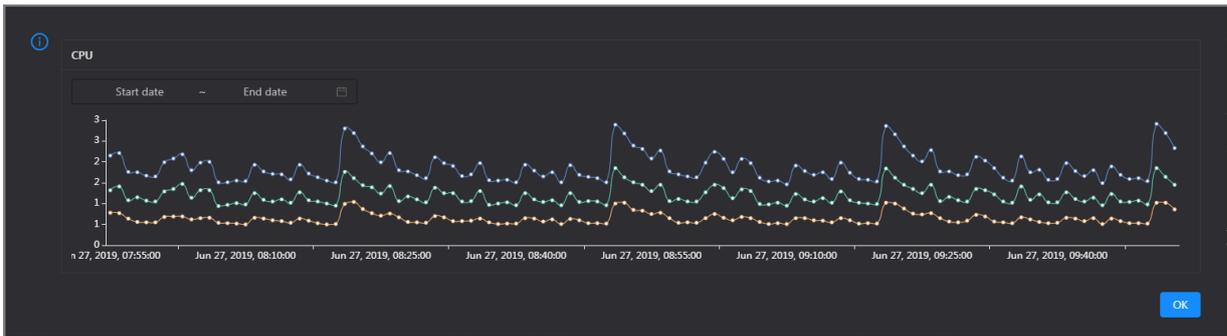
Service Role Status

This section displays the information about the services deployed on the host, including the roles, statuses, and number of services.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

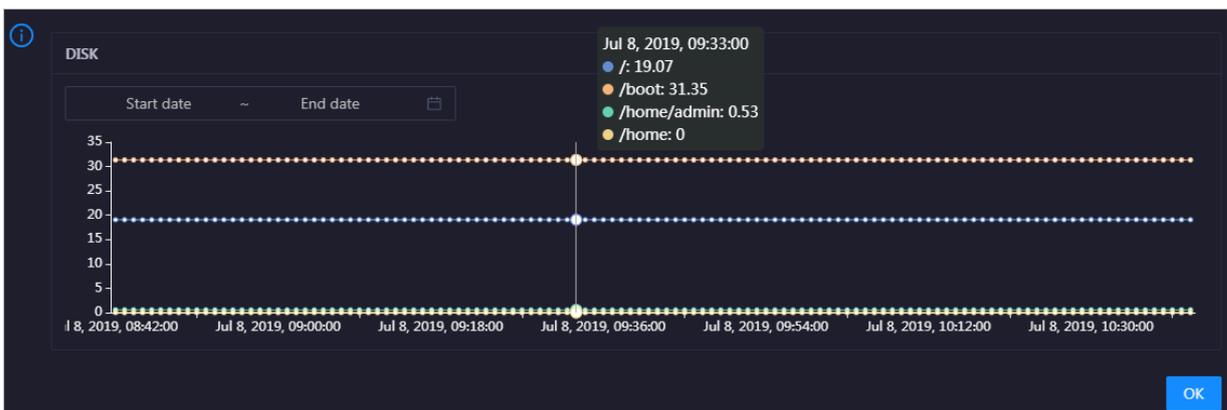


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

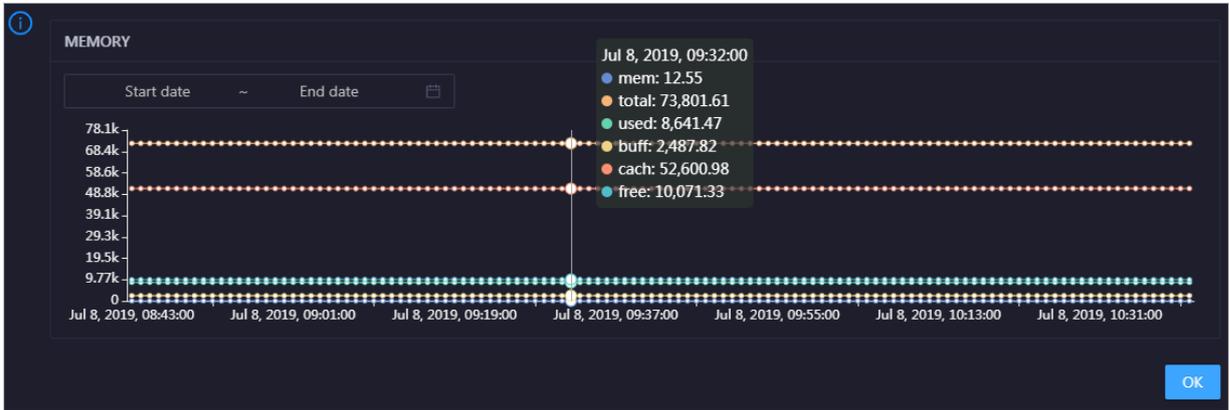


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (`mem`), total memory size (`total`), used memory size (`used`), size of memory used by kernel buffers (`buff`), size of memory used by the page cache (`cach`), and available memory size (`free`) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

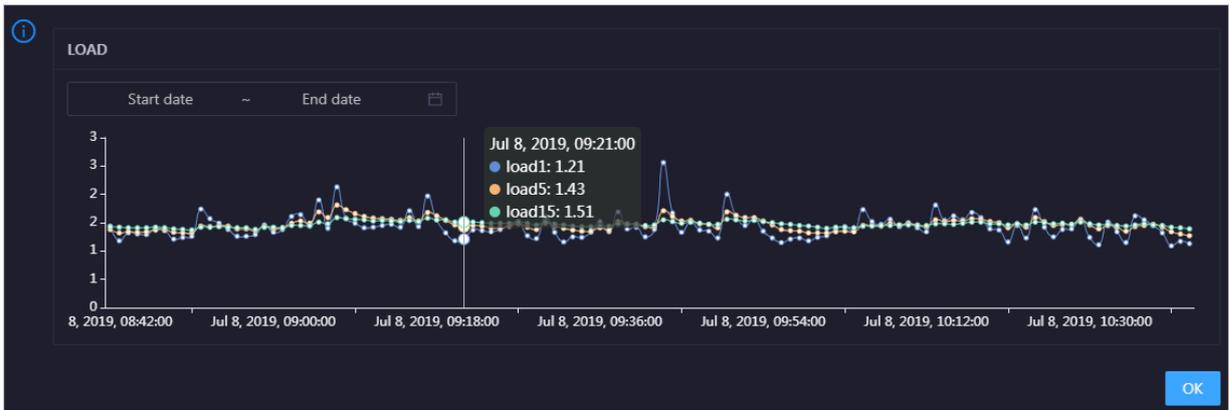


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

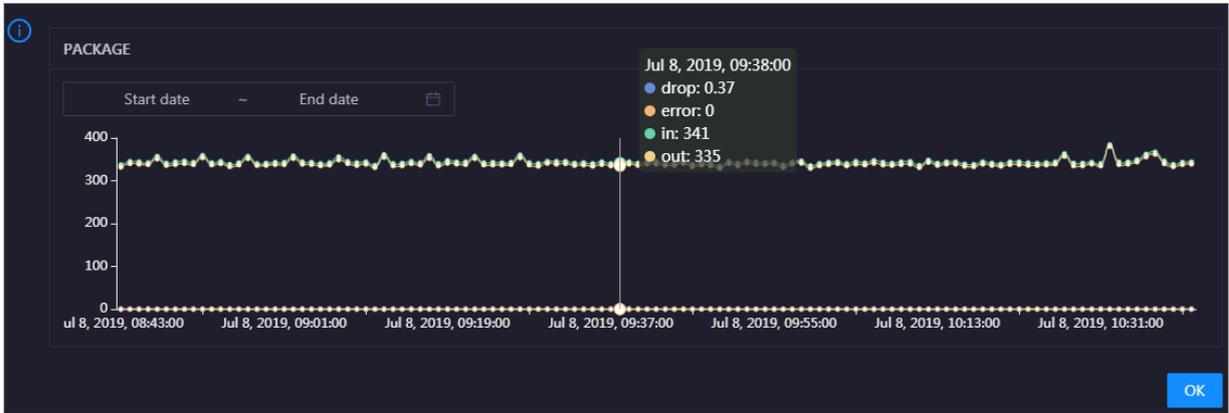


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

Health Check
View Details

Currently, 9 checkers are deployed on the service. 2 critical, 0 exception, and 0 warning alerts are reported.

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

Health Check History

This section displays a record of the health checks performed on the host.

Health Check History
View Details

Time	Event Content
Recently	1 alerts are reported by checkers.

1

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

You can click the event content of a check to view the exception items.

Checker
Host
Status
Status Updated At

bcc_host_live_check		CRITICAL	Jul 7, 2019, 18:35:30
---------------------	--	----------	-----------------------

1

3.1.8.5.2 Host charts

On the host chart page, you can view the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.

On the Hosts page, select a host in the left-side navigation pane, and then click the Charts tab. The Charts page for the host appears.



The Charts page displays trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host. For more information, see [Host overview](#).

3.1.8.5.3 Host health

On the host health status page, you can view the checkers of the selected host, including the checker details, check results, check history, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

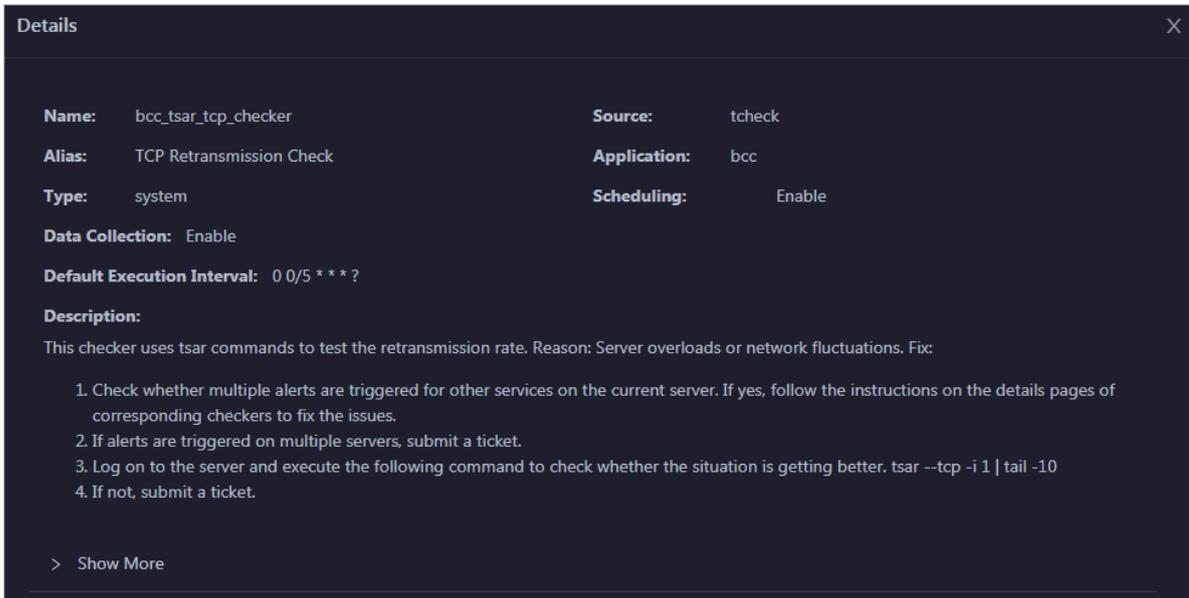
Entry

At the top of the O&M page, click the Hosts tab. On the page that appears, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

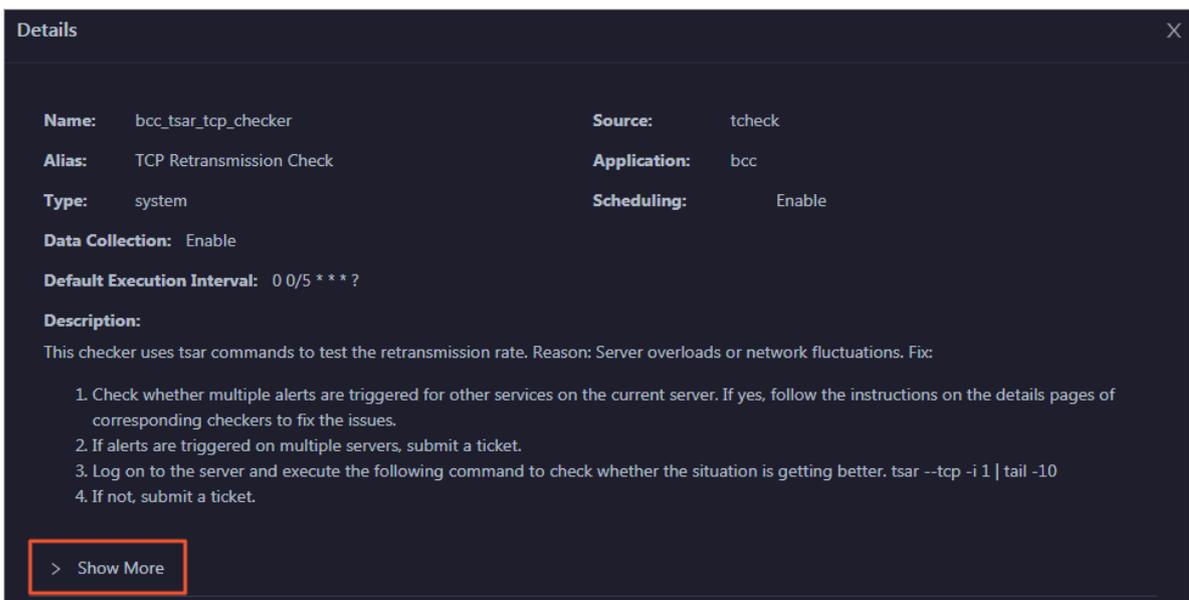
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

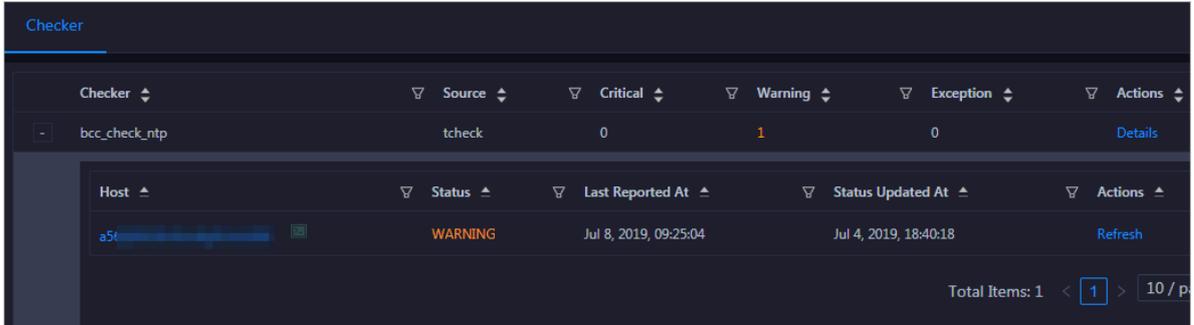


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

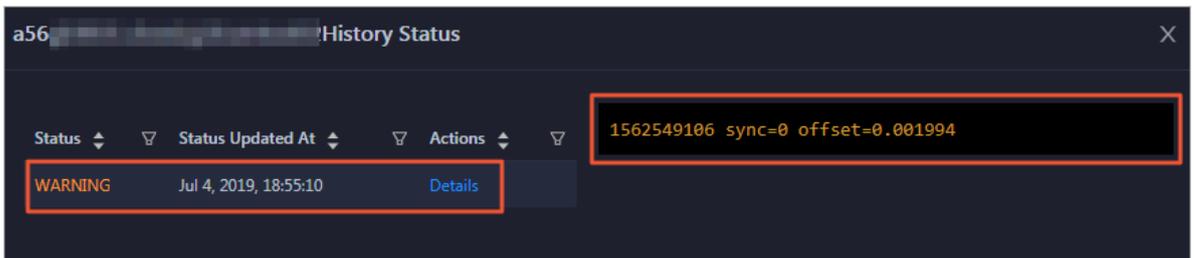
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

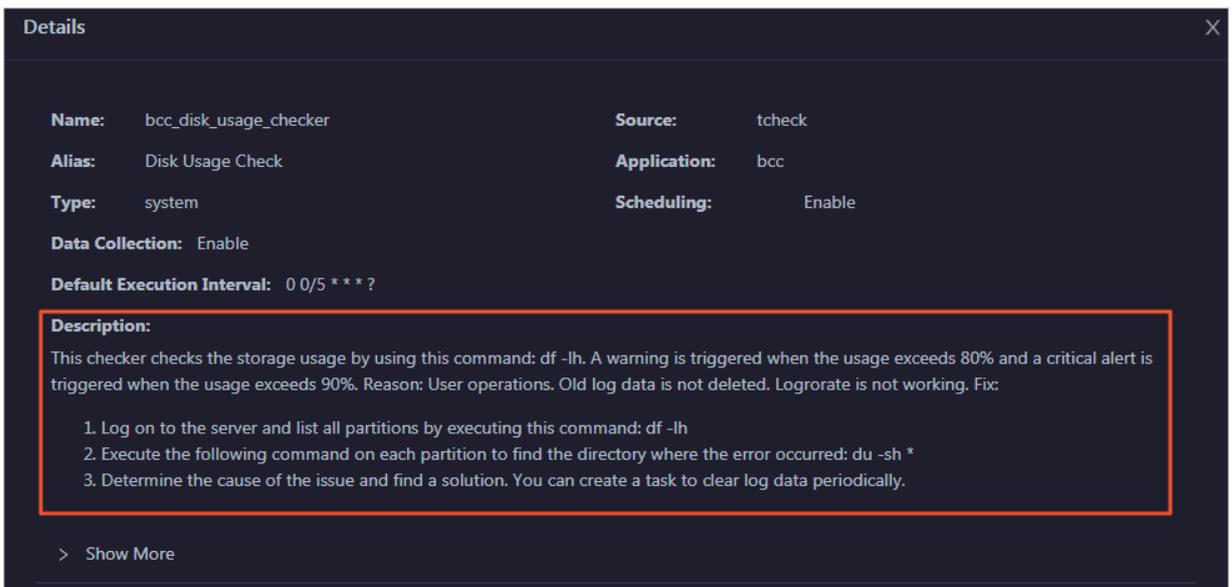


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

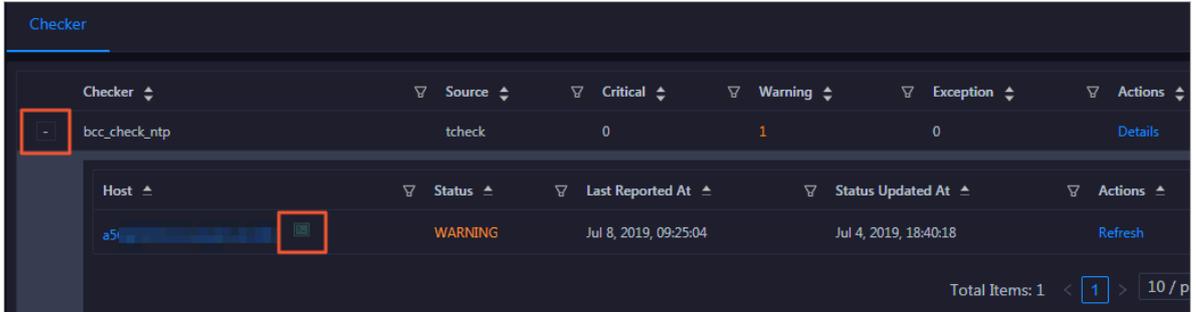
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



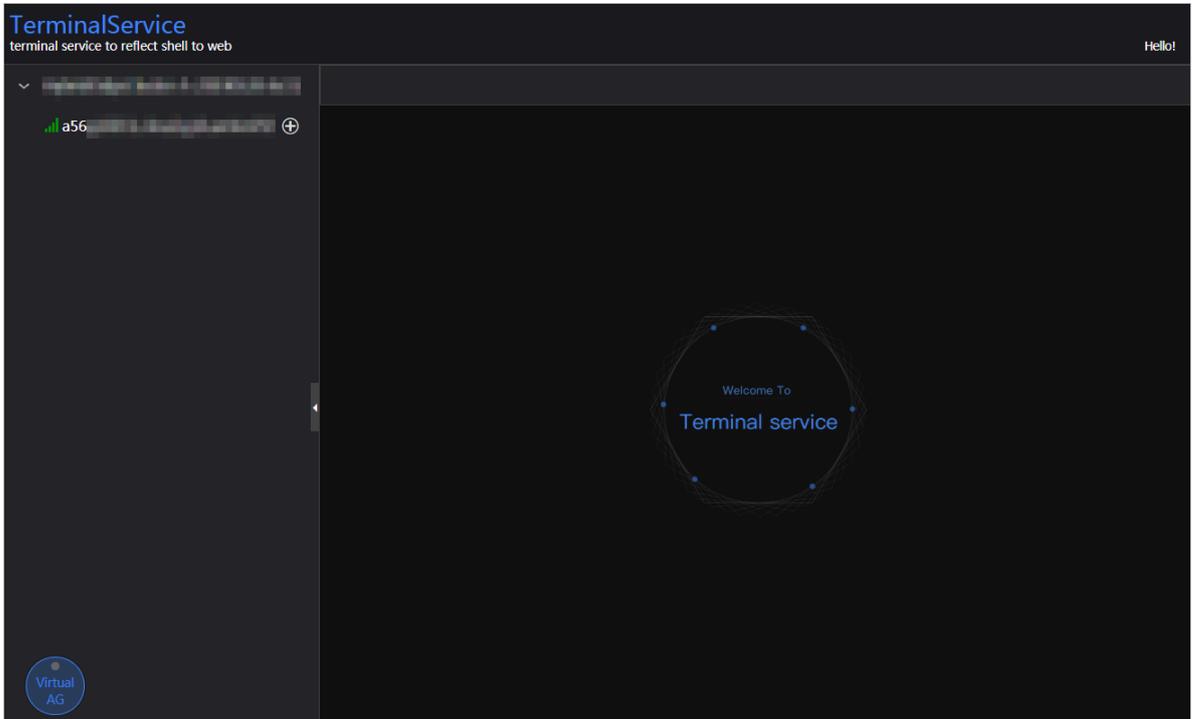
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

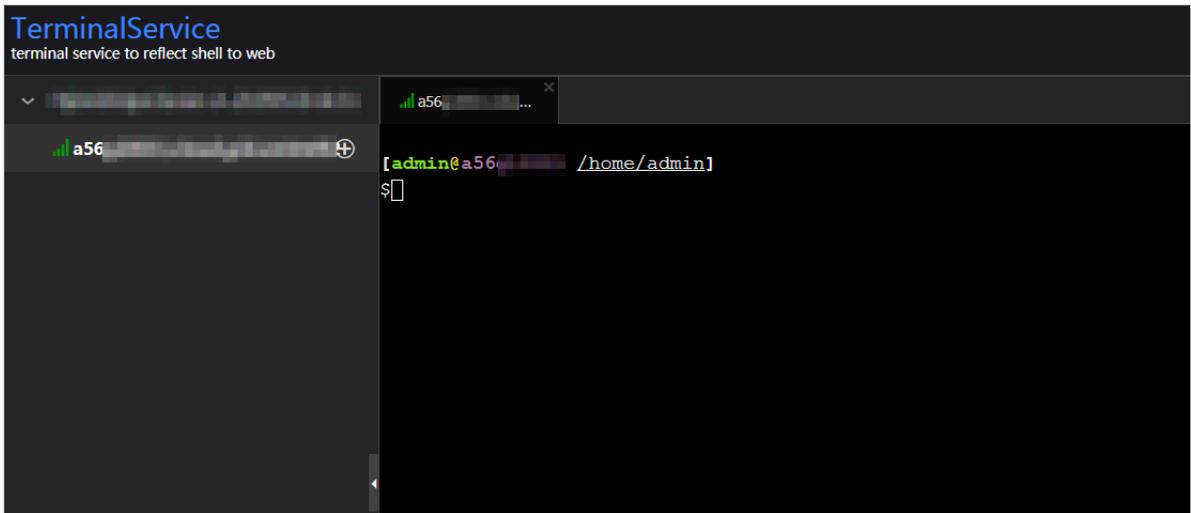
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

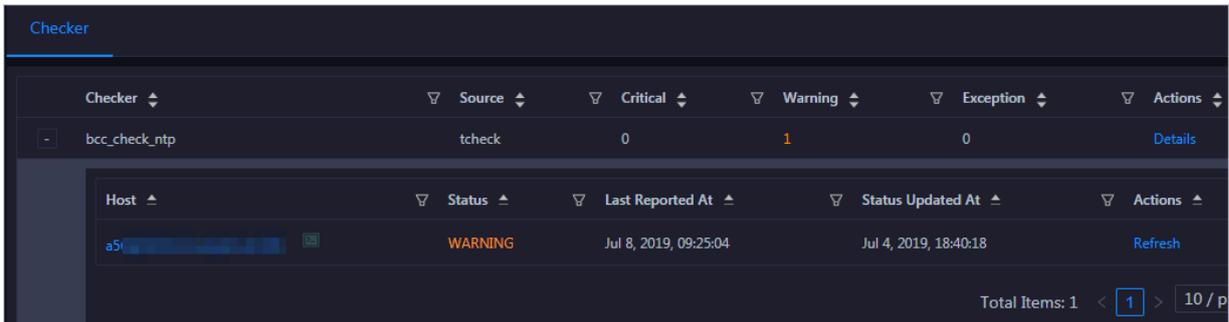


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.8.5.4 Host services

On the Services page, you can view information about service instances and service instance roles of a host.

On the Hosts page, select a host in the left-side navigation pane, and then click the Services tab. The Services page for the host appears.

On the Services page, you can view the cluster, service instances, and service instance roles of the host.

3.1.9 Dataphin

3.1.9.1 O&M overview

This topic describes the features of Dataphin O&M and how to access the Dataphin O&M page.

Modules

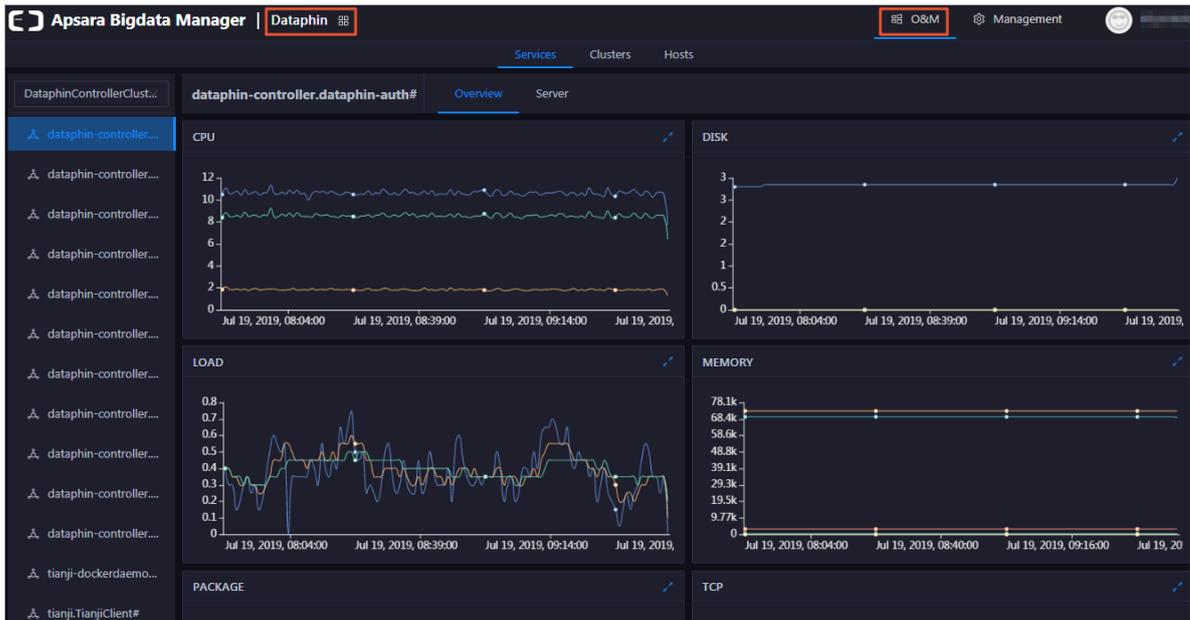
Dataphin O&M includes service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Service O&M	Service overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service in a cluster.
	Service hosts	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.
Cluster O&M	Cluster overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Cluster health	Displays the check results for a cluster. The check results are divided into the Critical, Warning, Exception, and OK types.
Host O&M	Host overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Host health	Displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click Dataphin.

3. On the page that appears, click O&M at the top. The Services page appears.



The O&M page includes three modules, namely, Services, Clusters, and Hosts.

3.1.9.2 Service O&M

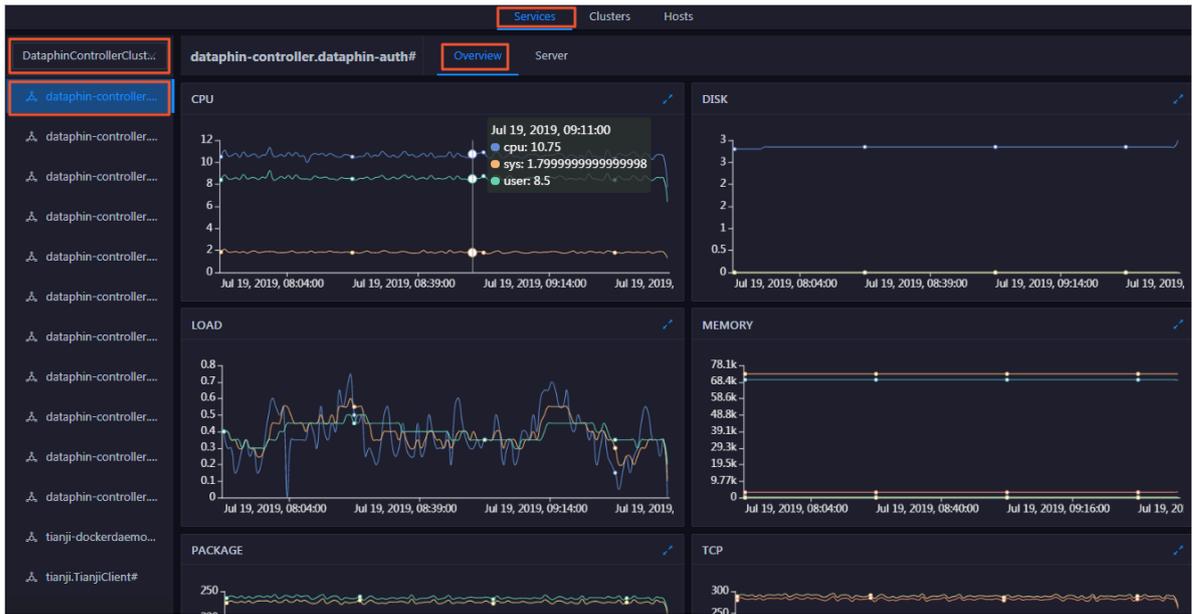
3.1.9.2.1 Service overview

The service overview page lists all Dataphin services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.

Entry

1. At the top of the O&M page, click Services.
2. On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.

3. Click the Overview tab. The Overview page for the service appears.



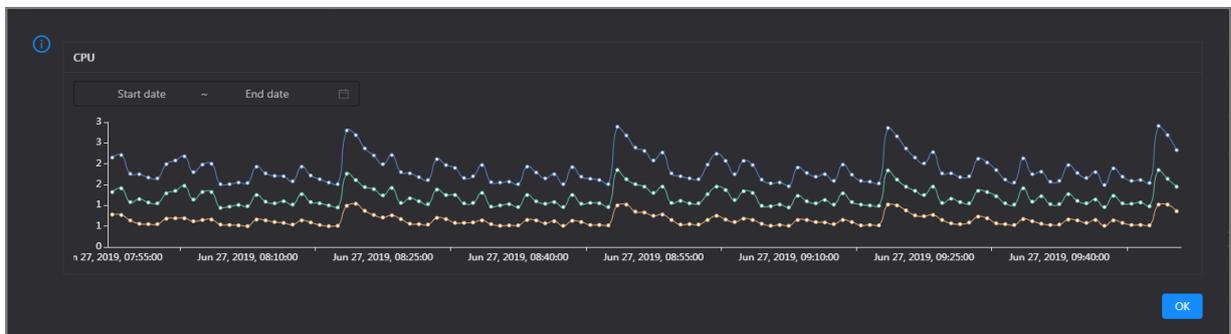
On the Overview page, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

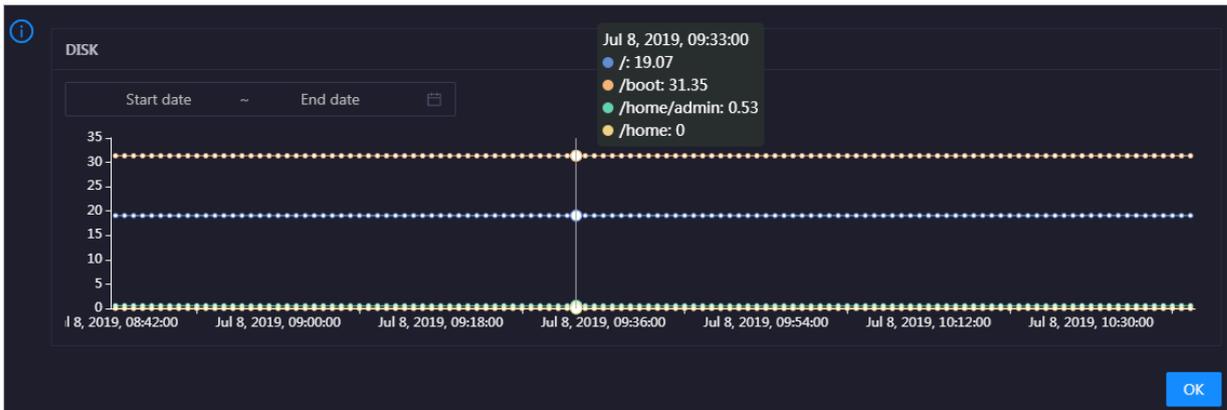
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the service in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the /, /boot, /home/admin, and /home directories for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

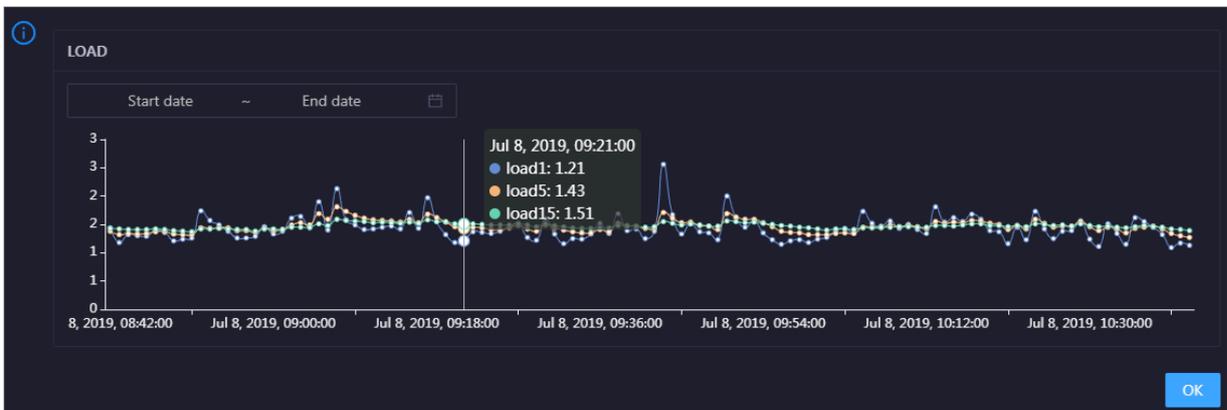


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

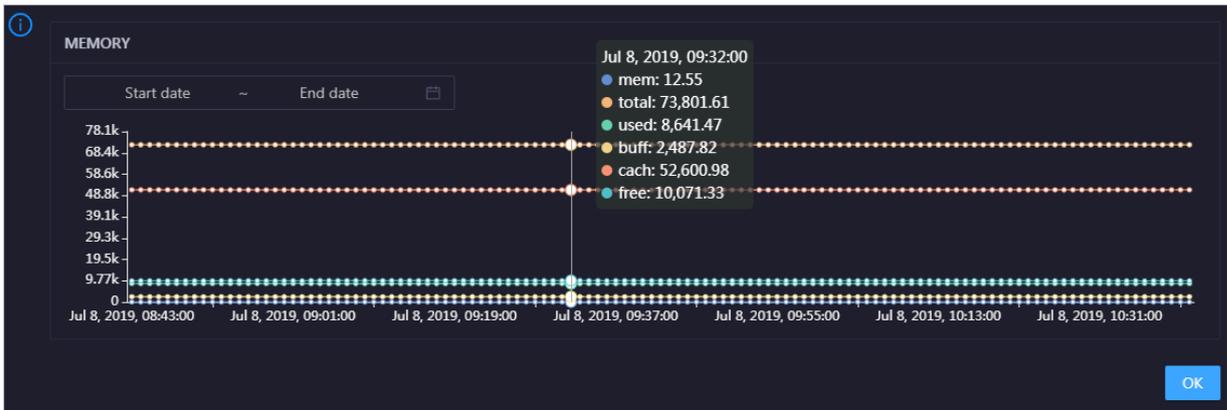


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

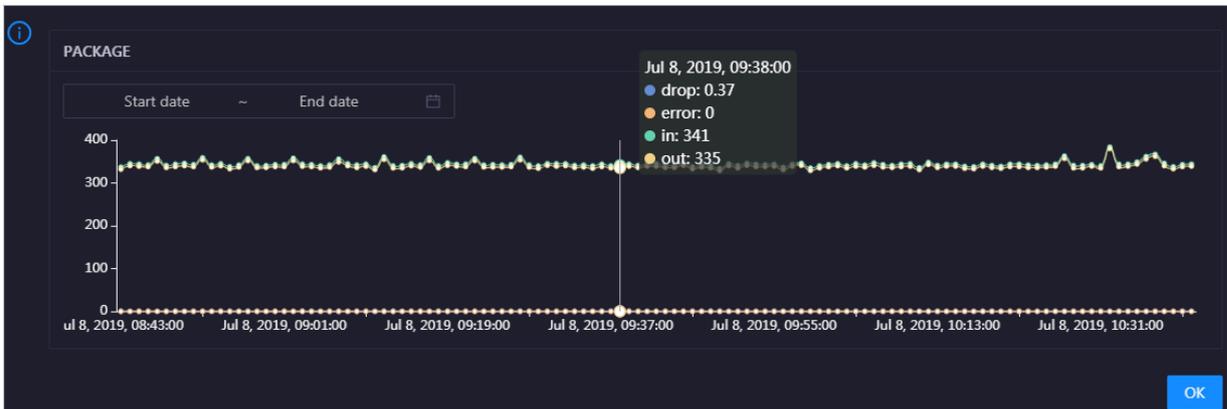


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the selected service over time in different colors. These trend lines reflect the data transmission status of the service.

Click  in the upper-right corner of the chart to zoom in it.



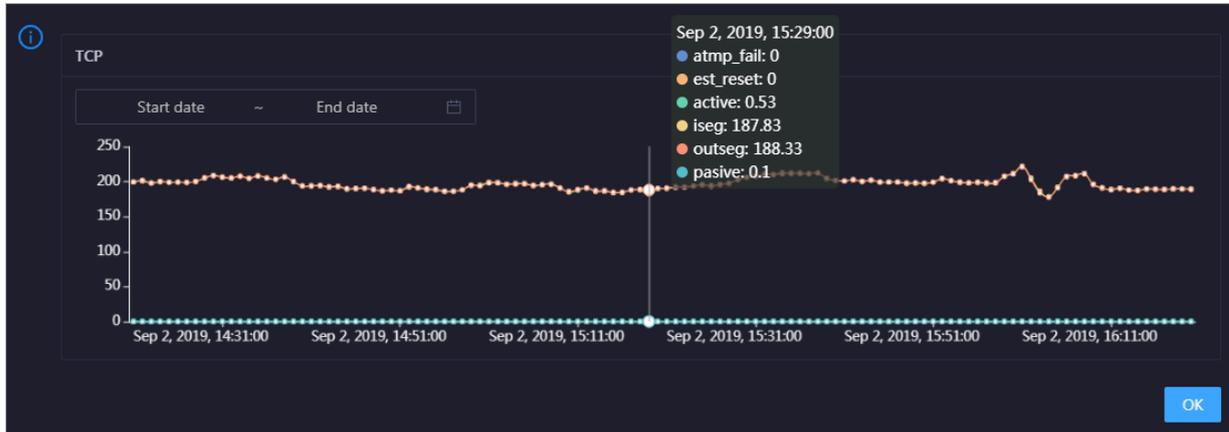
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP

connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the selected service over time in different colors. These trend lines reflect the TCP connection status of the service.

Click  in the upper-right corner of the chart to zoom in it.

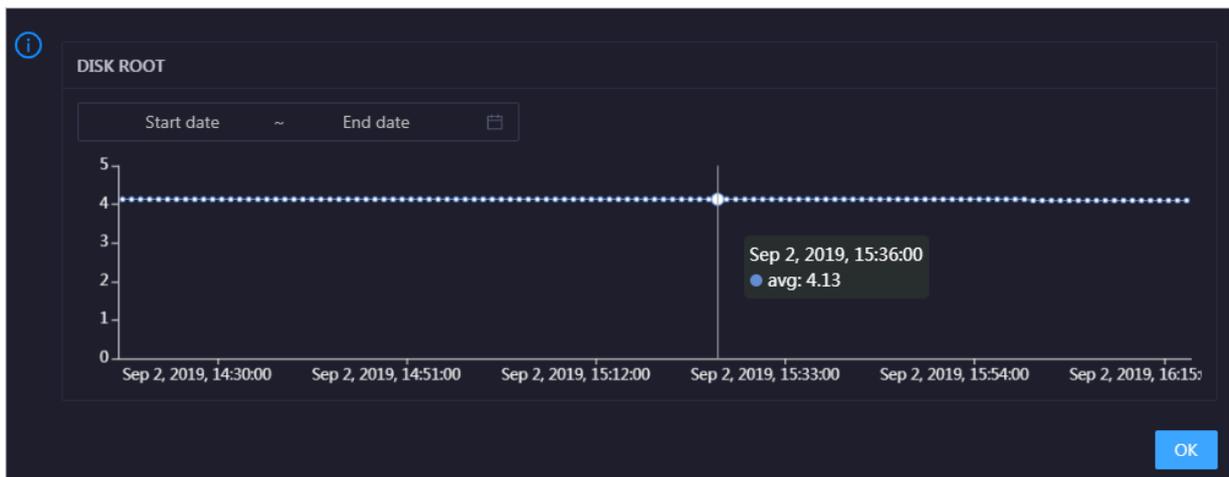


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

This chart displays the trend line of the average root disk usage (avg) for the selected service over time.

Click  in the upper-right corner of the chart to zoom in it.

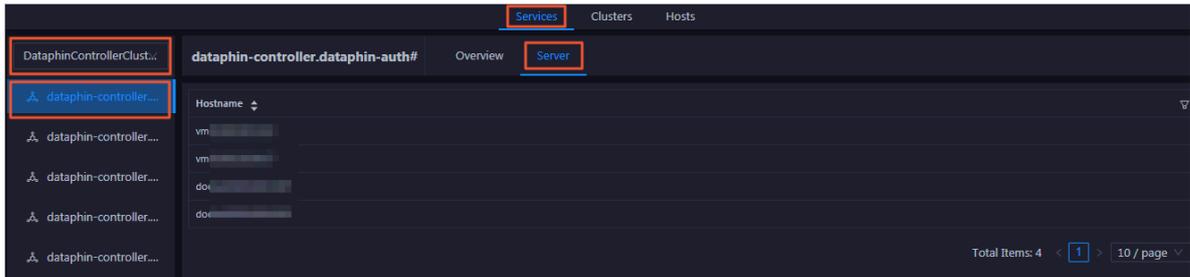


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

3.1.9.2.2 Service hosts

Apsara Bigdata Manager (ABM) allows you to view the host list of each Dataphin service so that you can understand the service deployment on hosts.

1. At the top of the O&M page, click **Services**.
2. On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.
3. Click the **Server** tab. The Server page for the service appears.



On the Server page, you can view the hosts where the selected service is run.

3.1.9.3 Cluster O&M

3.1.9.3.1 Cluster overview

The cluster overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.

Entry

1. At the top of the O&M page, click **Clusters**.

2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.



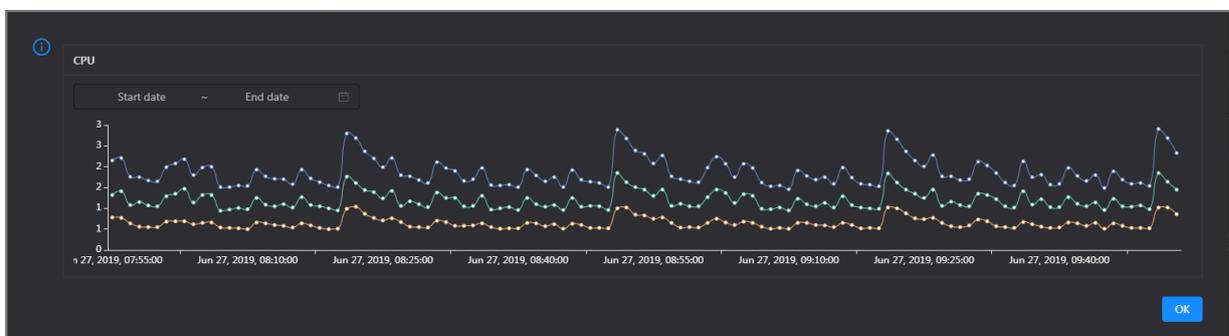
The Overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the cluster. To view information about a cluster, select a region in the left-side navigation pane, and then select a cluster in the region.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

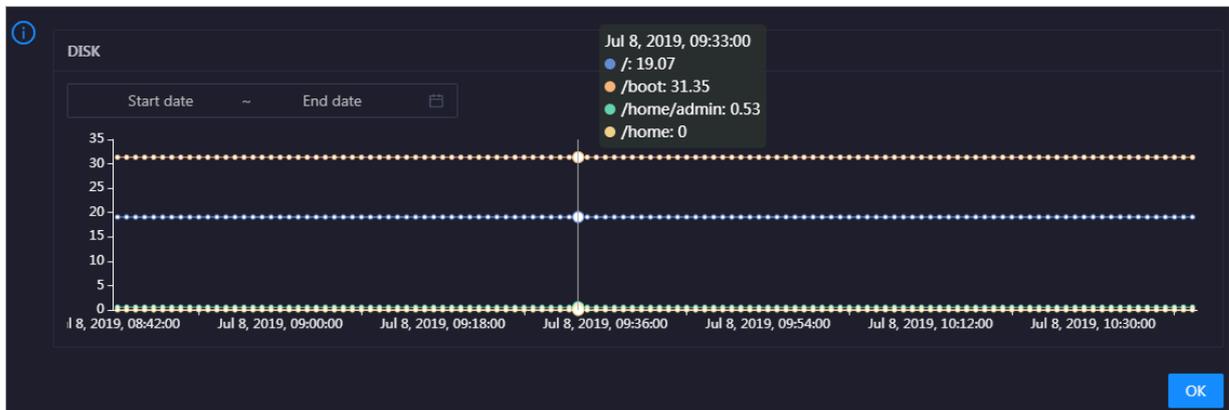
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

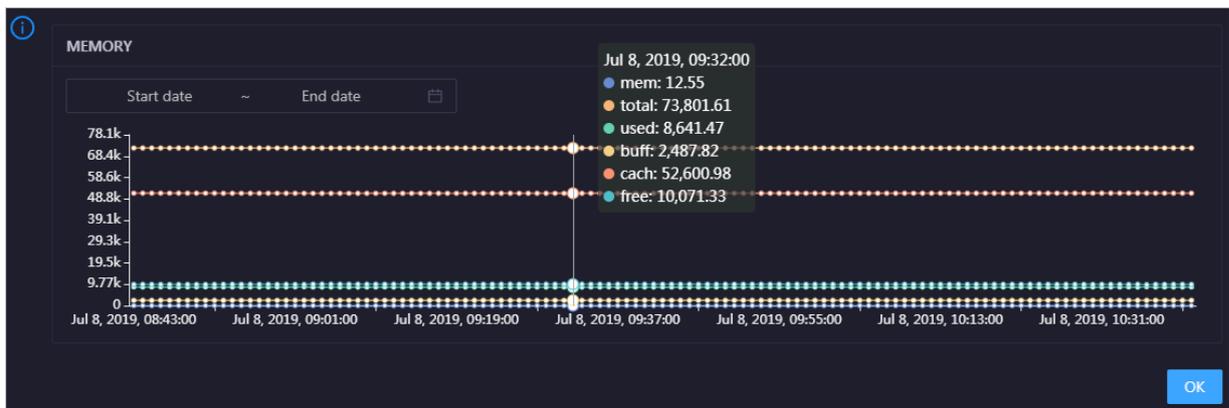


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

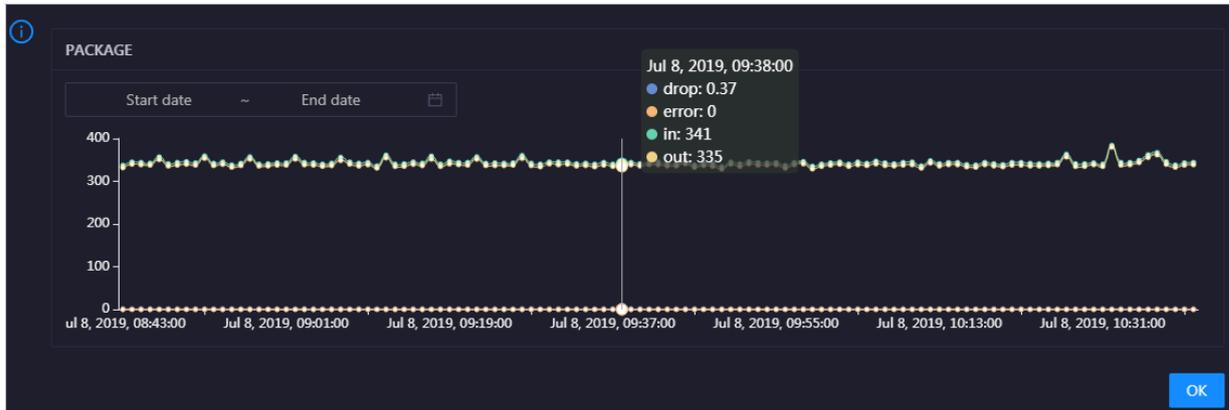


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

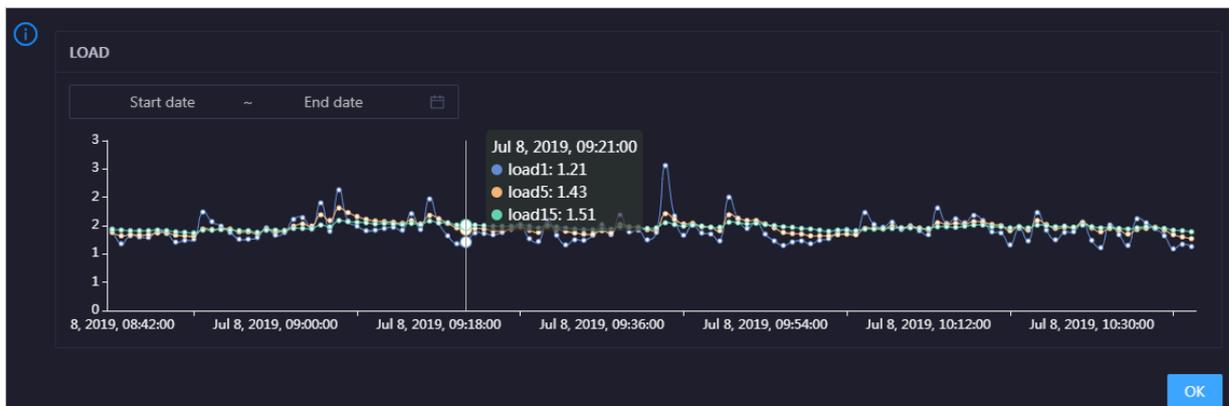


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

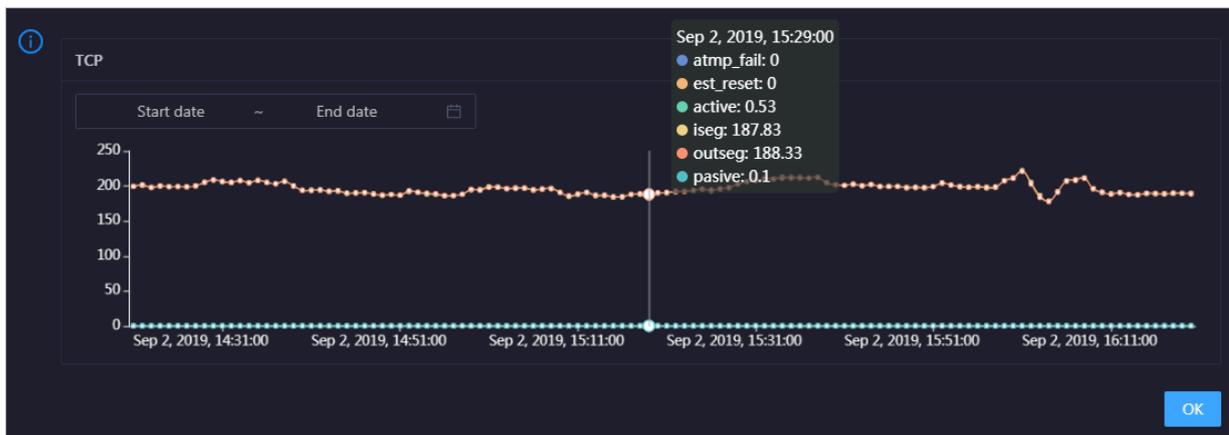


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

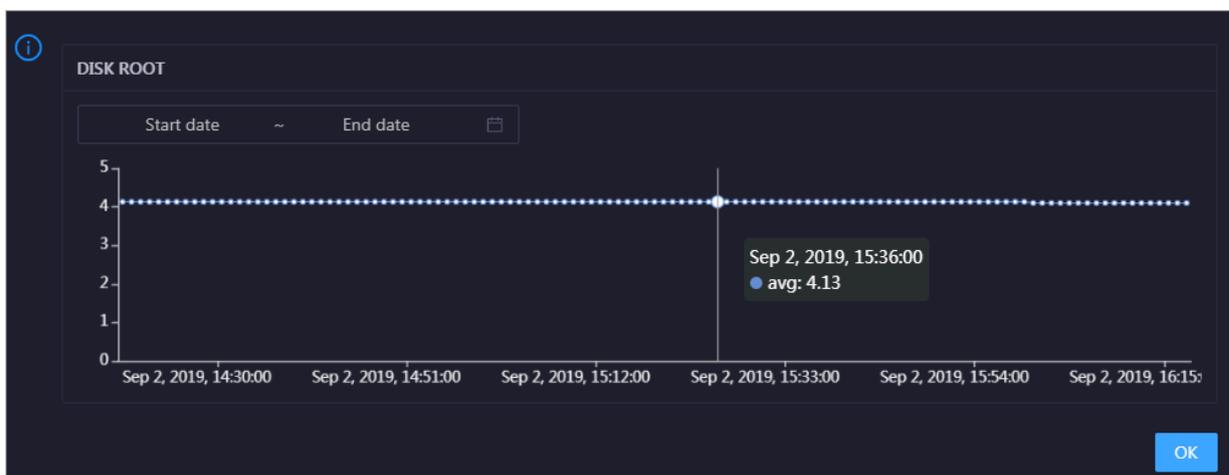


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

Click  in the upper-right corner of the chart to zoom in the chart.



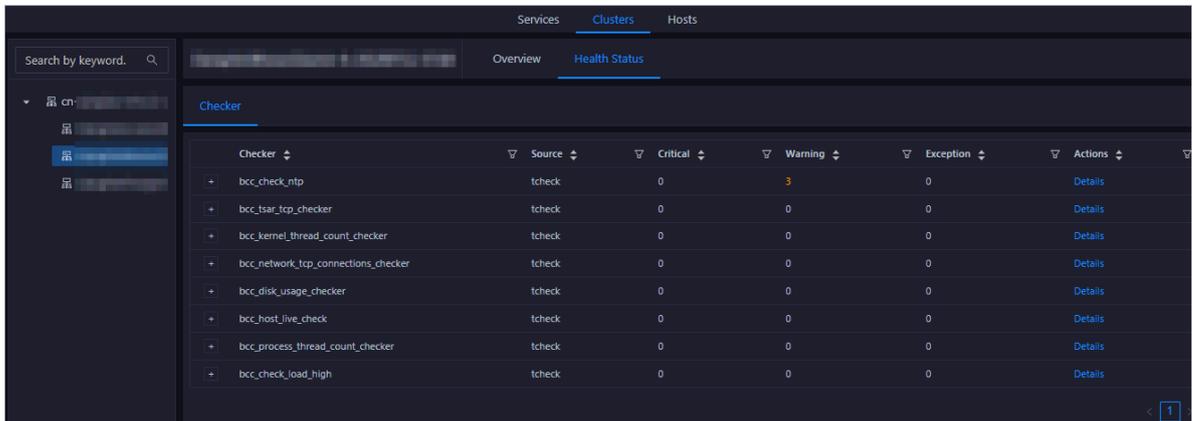
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

3.1.9.3.2 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. At the top of the O&M page, click Clusters.
2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.

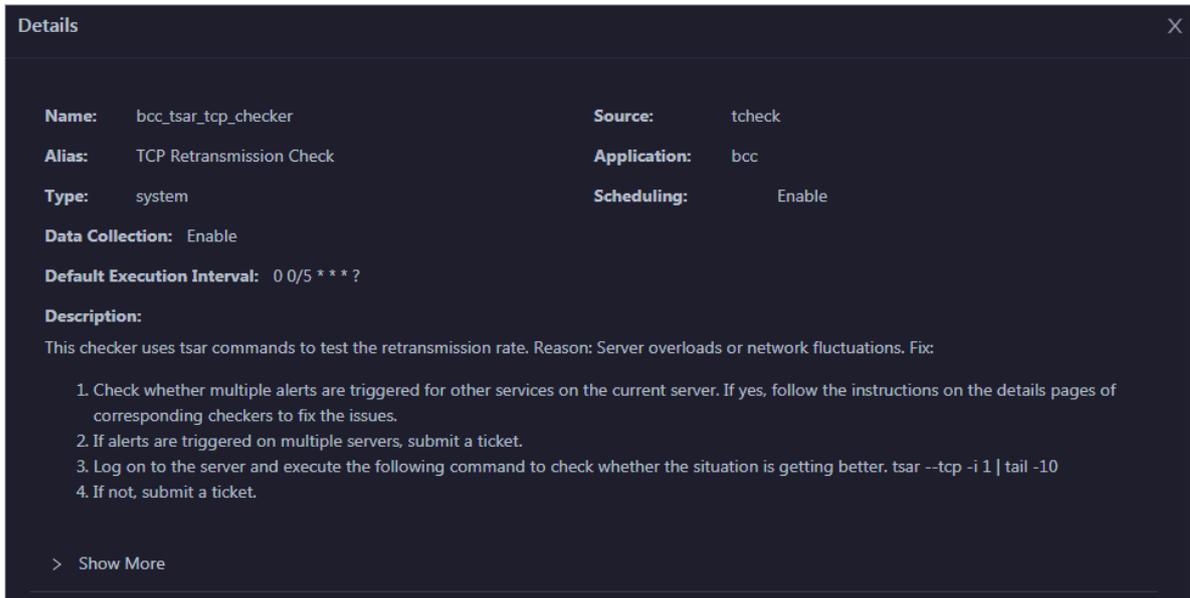


Checker	Source	Critical	Warning	Exception	Actions
+ bcc_check_ntp	tcheck	0	3	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_network_tcp_connections_checker	tcheck	0	0	0	Details
+ bcc_disk_usage_checker	tcheck	0	0	0	Details
+ bcc_host_live_check	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

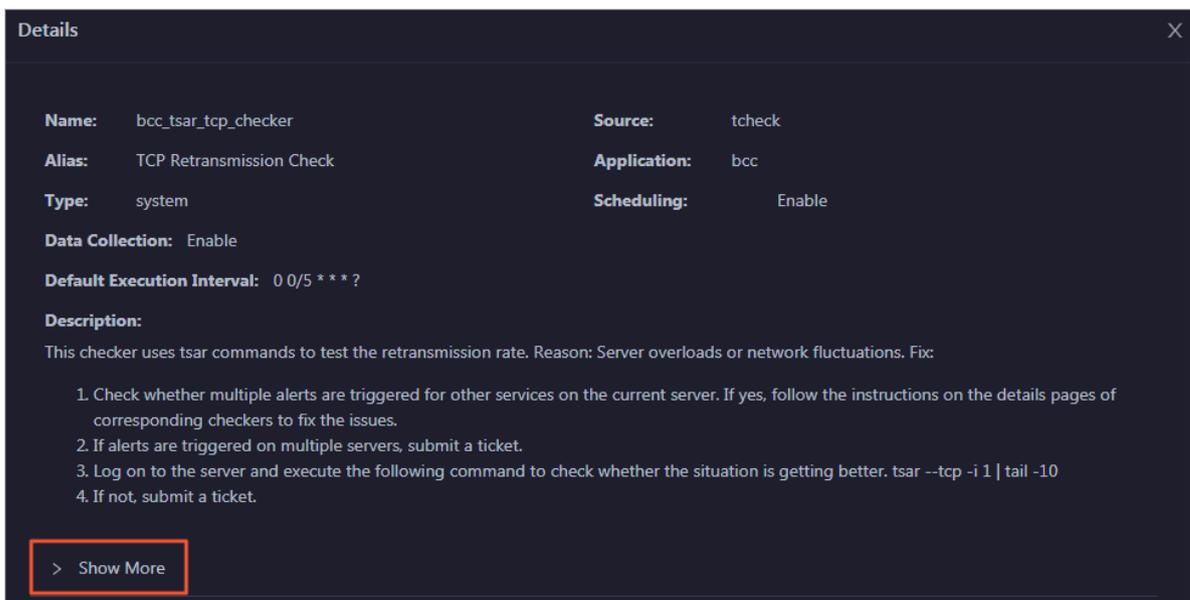
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

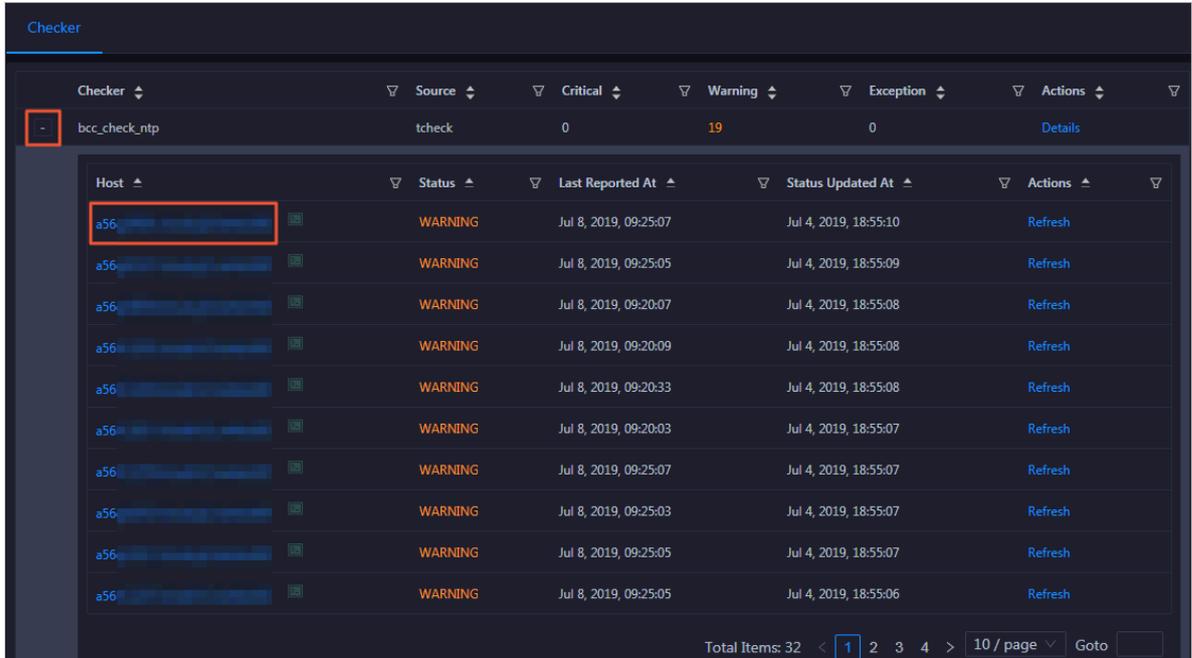


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

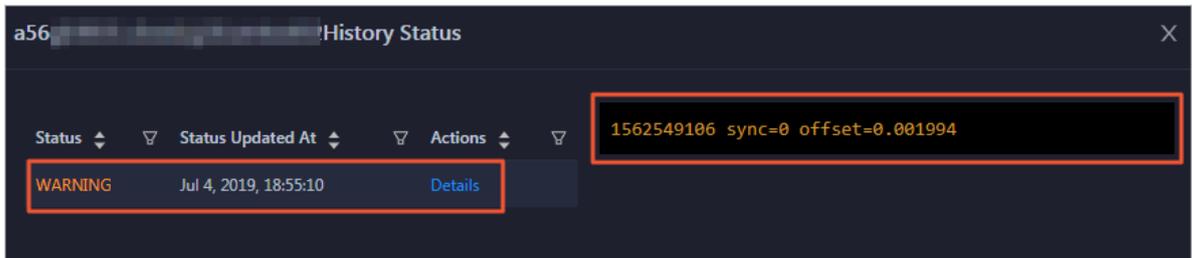
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

- 1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.**



- 2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**



Clear alerts

On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.

Details ✕

Name: bcc_disk_usage_checker	Source: tcheck
Alias: Disk Usage Check	Application: bcc
Type: system	Scheduling: Enable
Data Collection: Enable	
Default Execution Interval: 0 0/5 * * * ?	

Description:
 This checker checks the storage usage by using this command: `df -lh`. A warning is triggered when the usage exceeds 80% and a critical alert is triggered when the usage exceeds 90%. Reason: User operations. Old log data is not deleted. Logrotate is not working. Fix:

1. Log on to the server and list all partitions by executing this command: `df -lh`
2. Execute the following command on each partition to find the directory where the error occurred: `du -sh *`
3. Determine the cause of the issue and find a solution. You can create a task to clear log data periodically.

> Show More

Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

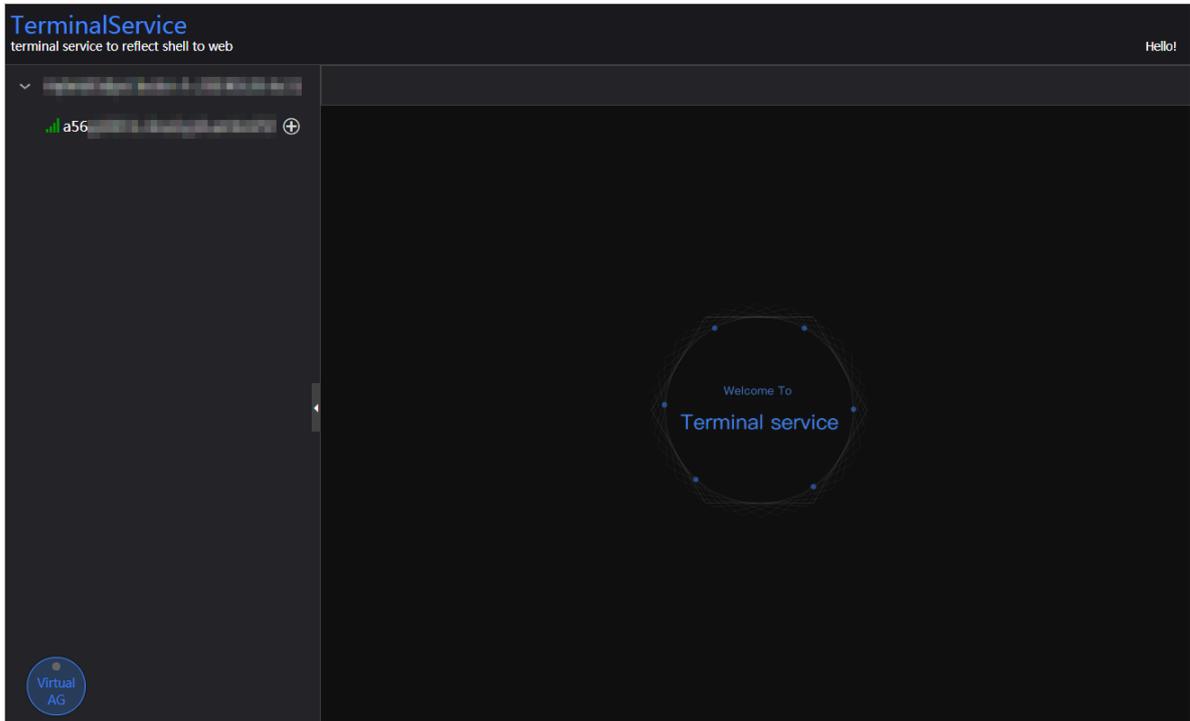
- 1. On the Health Status page, click + to expand a checker with alerts.**

Checker

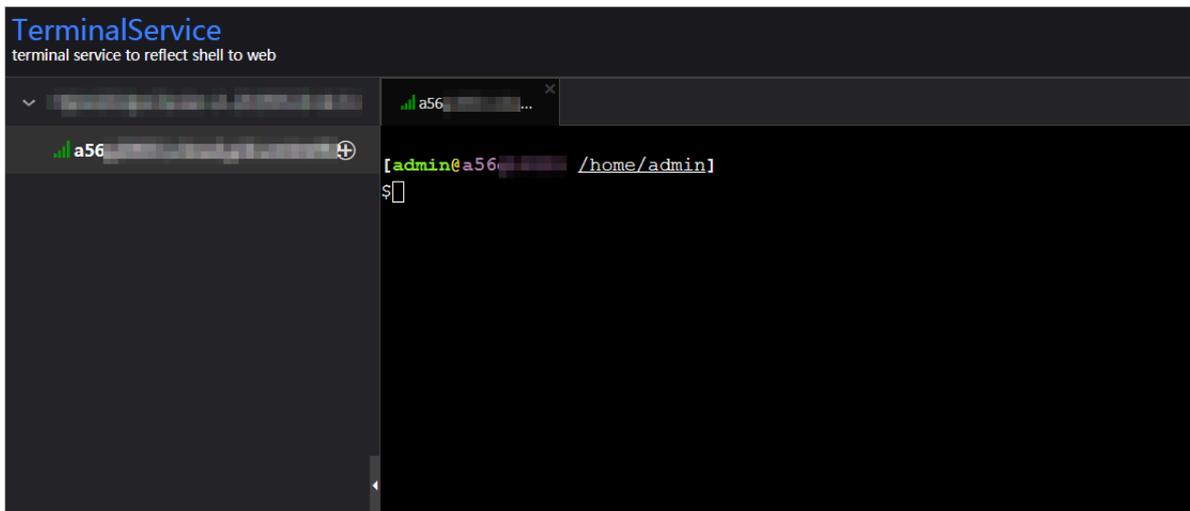
Checker	Source	Critical	Warning	Exception	Actions
- bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56 +	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56 +	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56 +	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56 +	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.



3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

The screenshot shows a 'Checker' interface with a table of hosts. The table has columns for Host, Status, Last Reported At, Status Updated At, and Actions. The first row shows a host with a 'WARNING' status and a 'Refresh' button highlighted in a red box.

Checker	Source	Critical	Warning	Exception	Actions
bcc_check_ntp	tcheck	0	19	0	Details
Host	Status	Last Reported At	Status Updated At	Actions	
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh	
a56	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh	
a56	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh	
a56	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh	
a56	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh	
a56	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh	
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh	

3.1.9.4 Host O&M

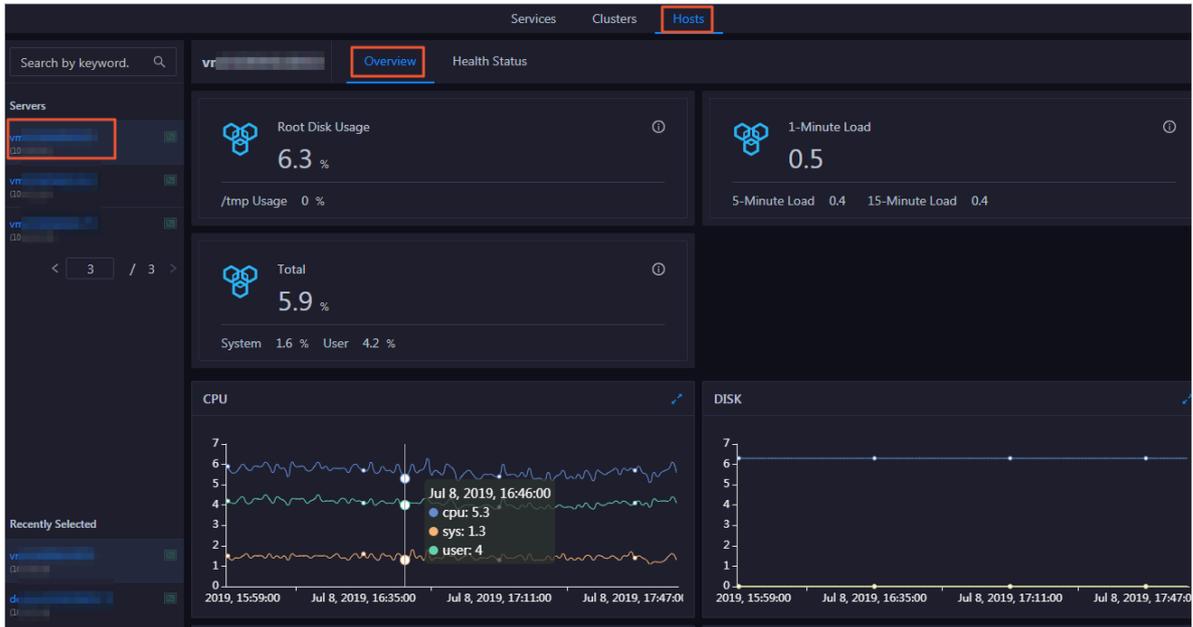
3.1.9.4.1 Host overview

The host overview page displays the overall running information about a host in a Dataphin cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Entry

1. At the top of the O&M page, click Hosts.

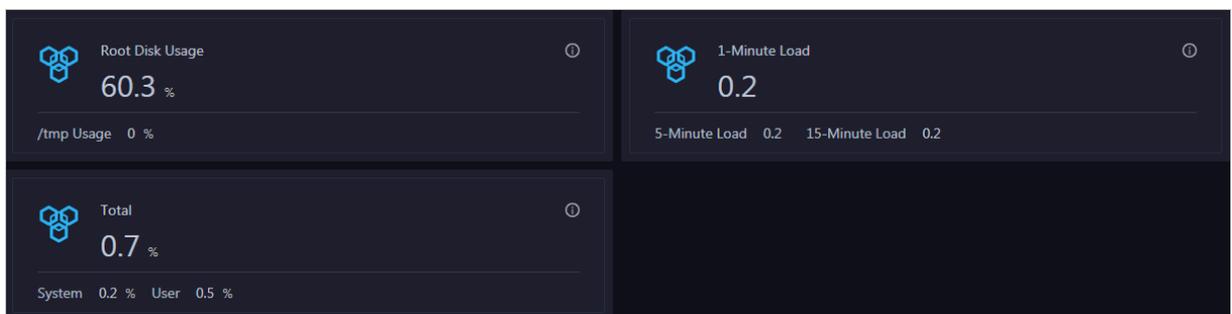
2. On the Hosts page, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.



On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Root Disk Usage, Total, and 1-Minute Load

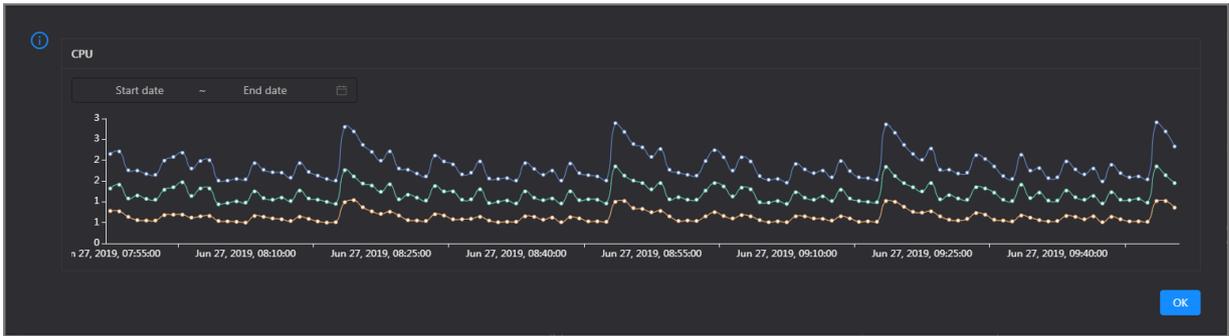
These sections display the root disk usage, total usage, and 1-minute load for the selected host. The Root Disk Usage section provides the usage of the `/tmp` directory. The Total section provides the system usage and user usage. The 1-Minute Load section provides the 1-minute, 5-minute, and 15-minute load averages.



CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

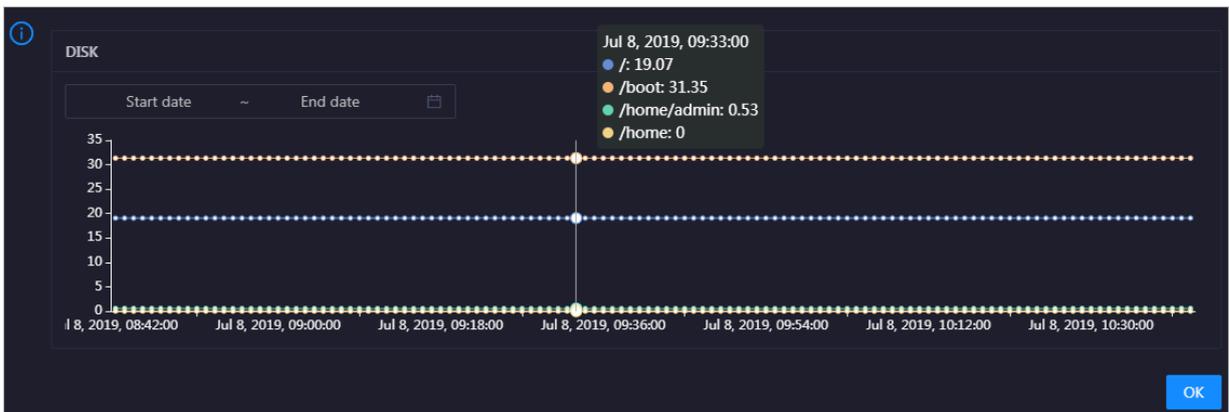


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

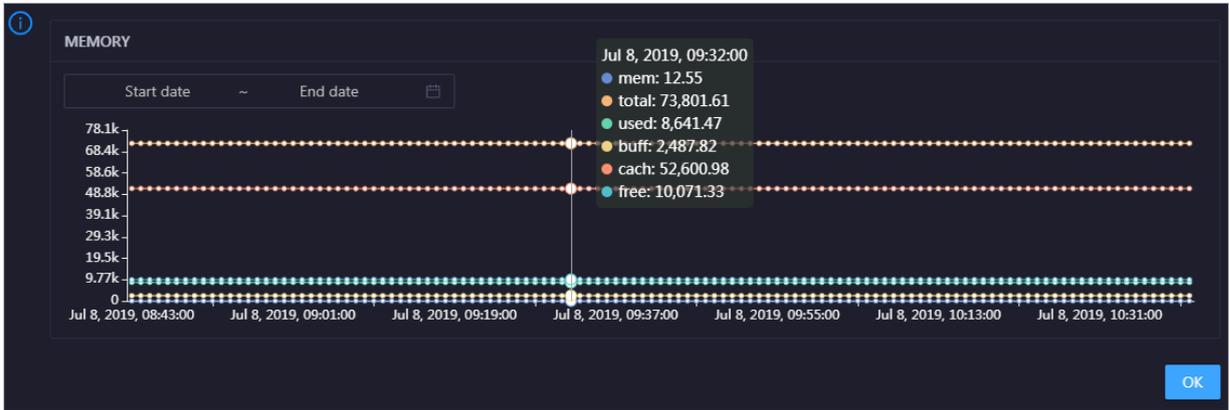


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (`mem`), total memory size (`total`), used memory size (`used`), size of memory used by kernel buffers (`buff`), size of memory used by the page cache (`cach`), and available memory size (`free`) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

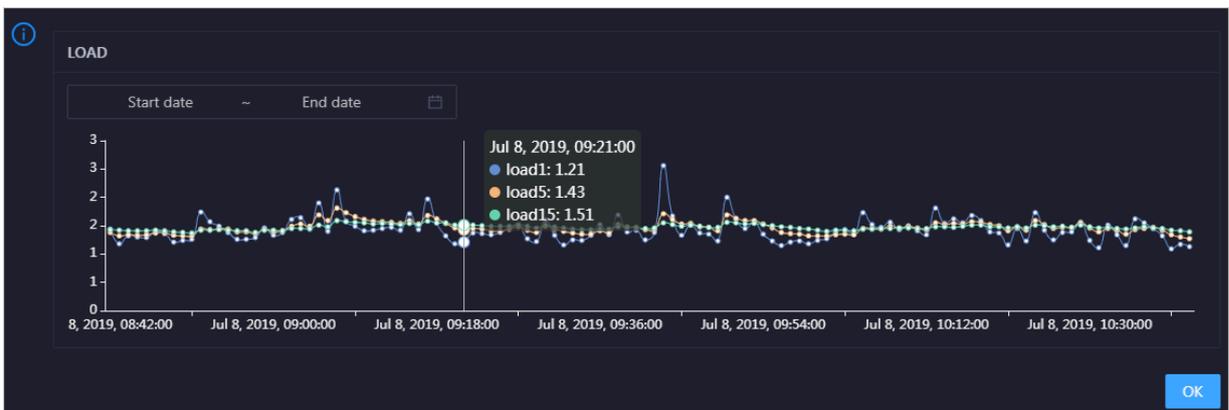


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.

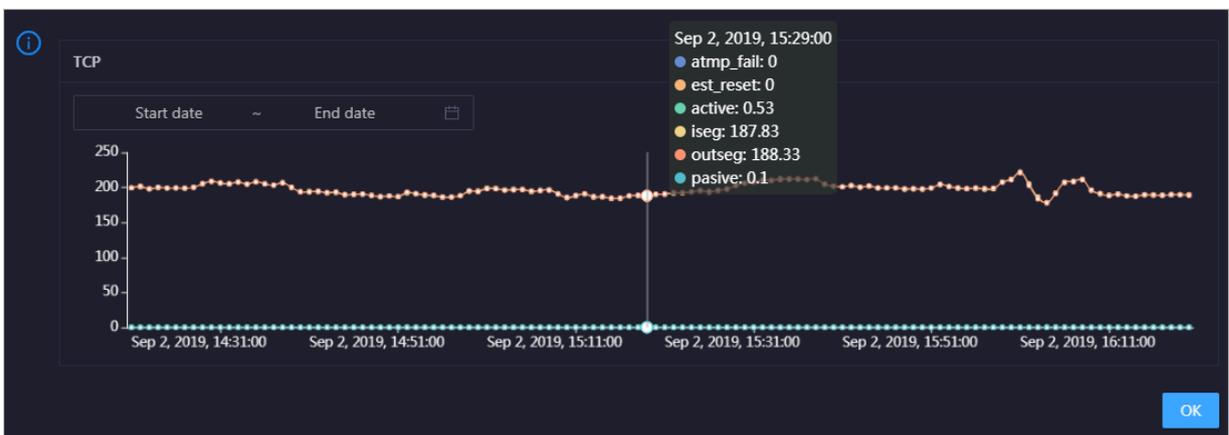


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click  in the upper-right corner of the chart to zoom in the chart.

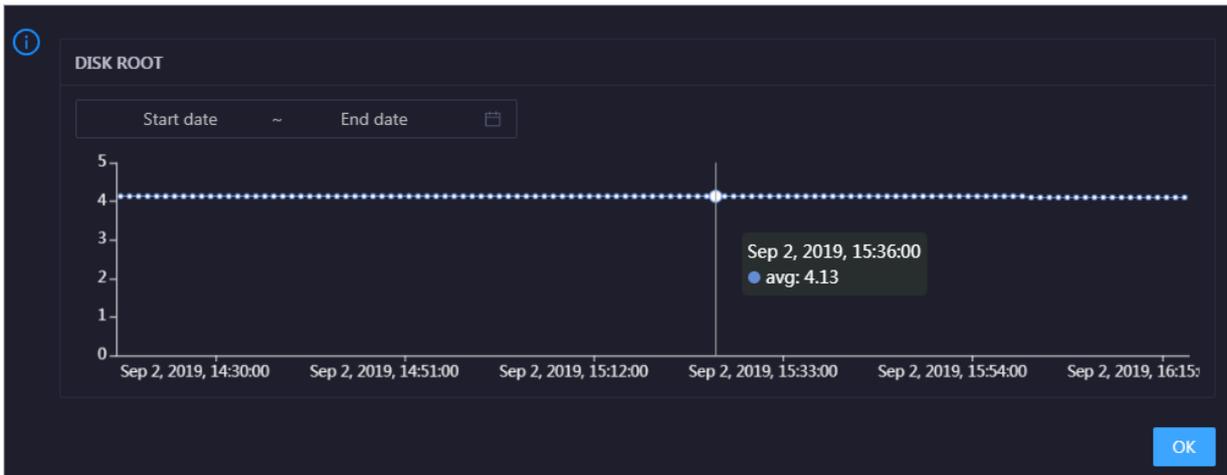


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the host over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

Health Check [View Details](#)

Currently, 9 checkers are deployed on the service. 2 critical, 0 exception, and 0 warning alerts are reported.

Click [View Details](#) to go to the [Host health](#) page. On this page, you can view the health check details.

Health Check History

This section displays a record of the health checks performed on the host.

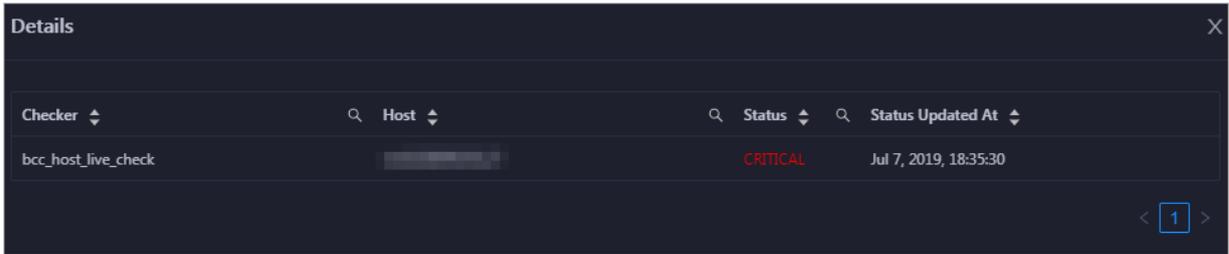
Health Check History [View Details](#)

Time	Event Content
Recently	1 alerts are reported by checkers.

< 1 >

Click [View Details](#) to go to the [Host health](#) page. On this page, you can view the health check details.

You can click the event content of a check to view the exception items.

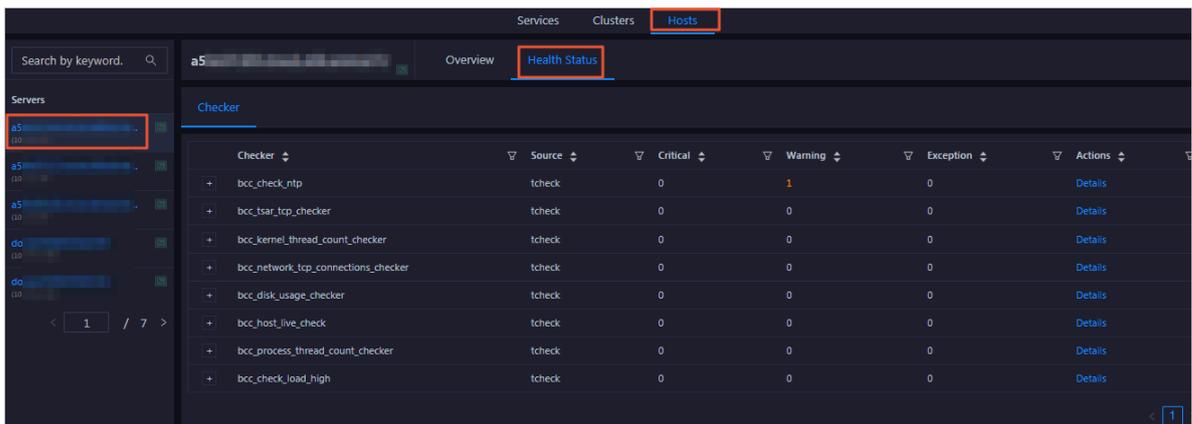


3.1.9.4.2 Host health

On the host health status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

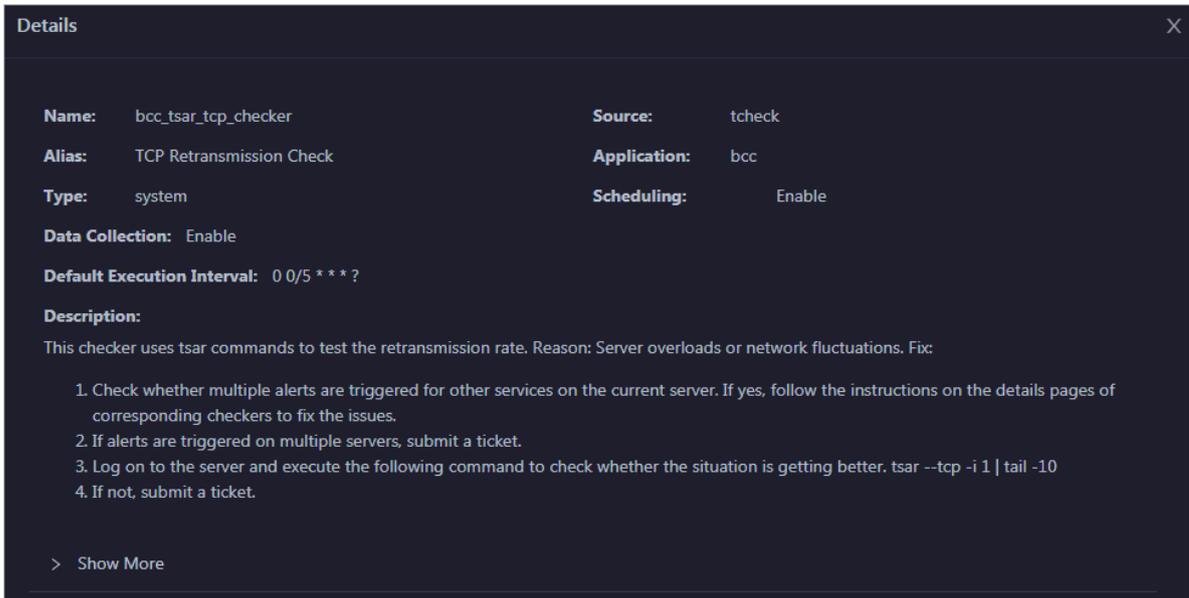
1. At the top of the O&M page, click the Hosts tab.
2. On the Hosts page that appears, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.



On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

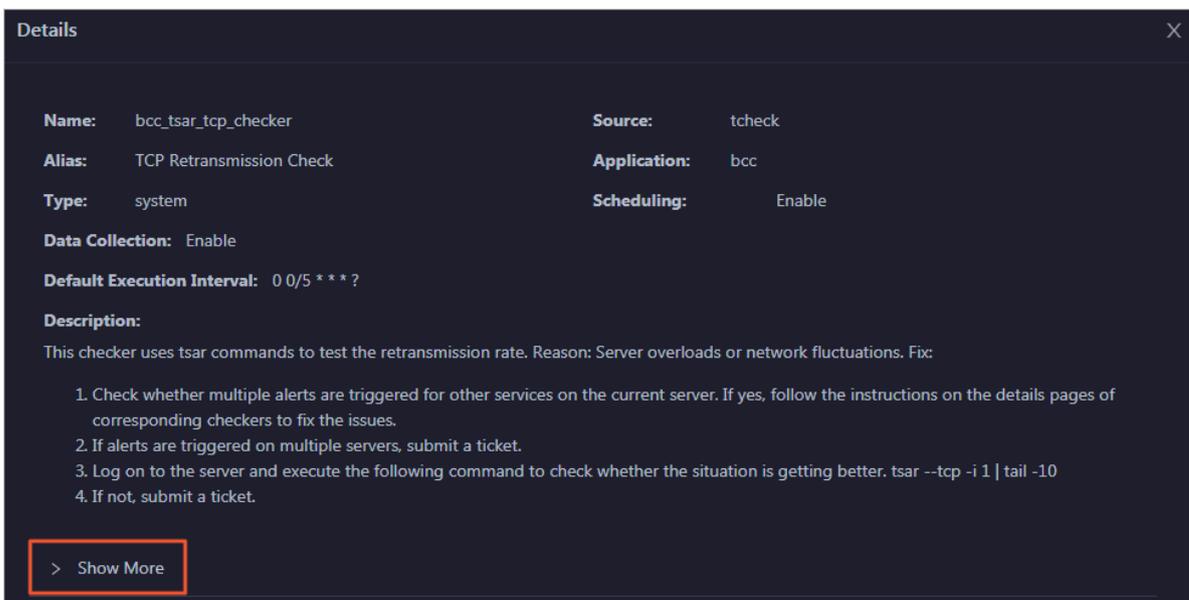
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

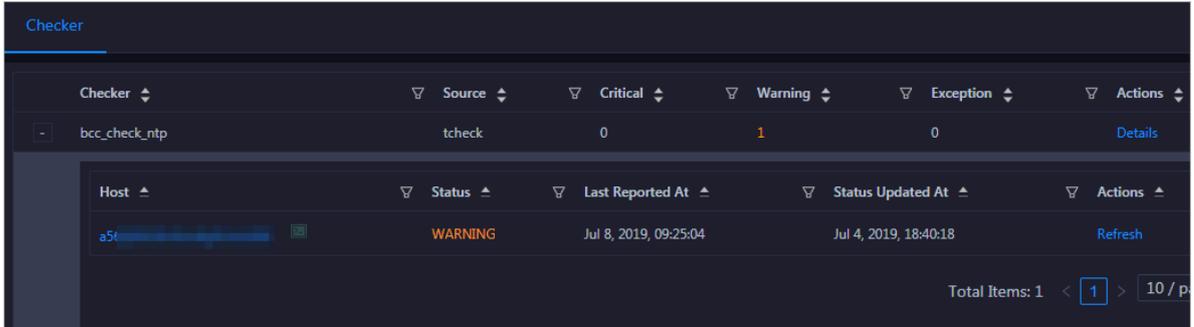


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

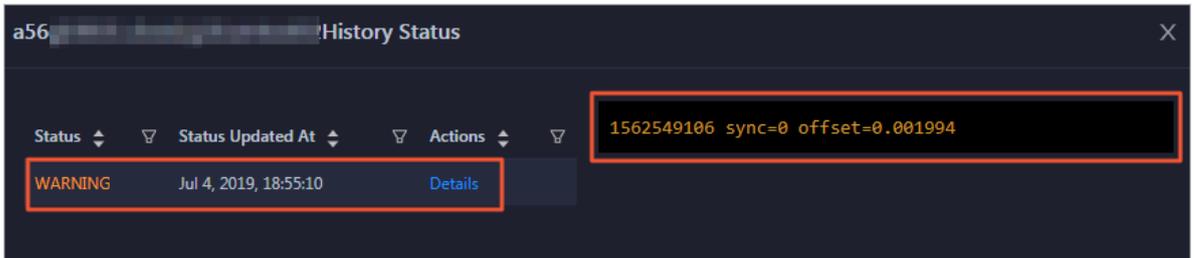
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

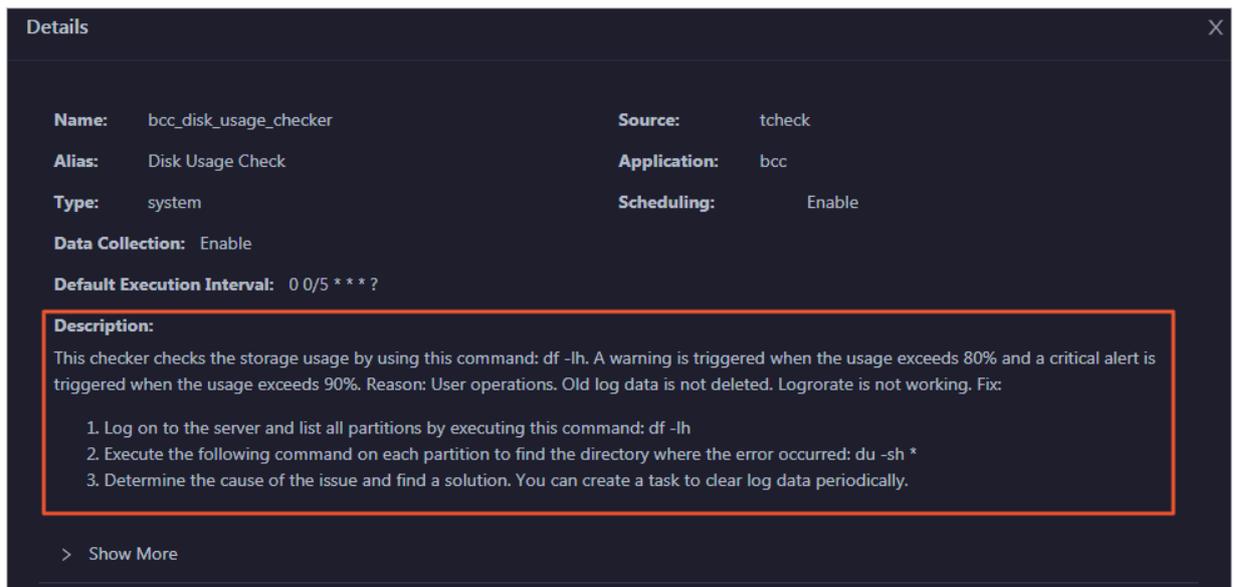


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

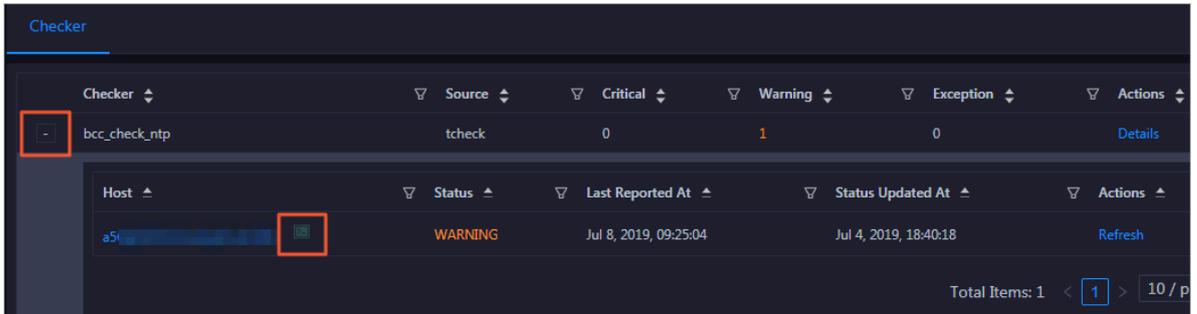
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



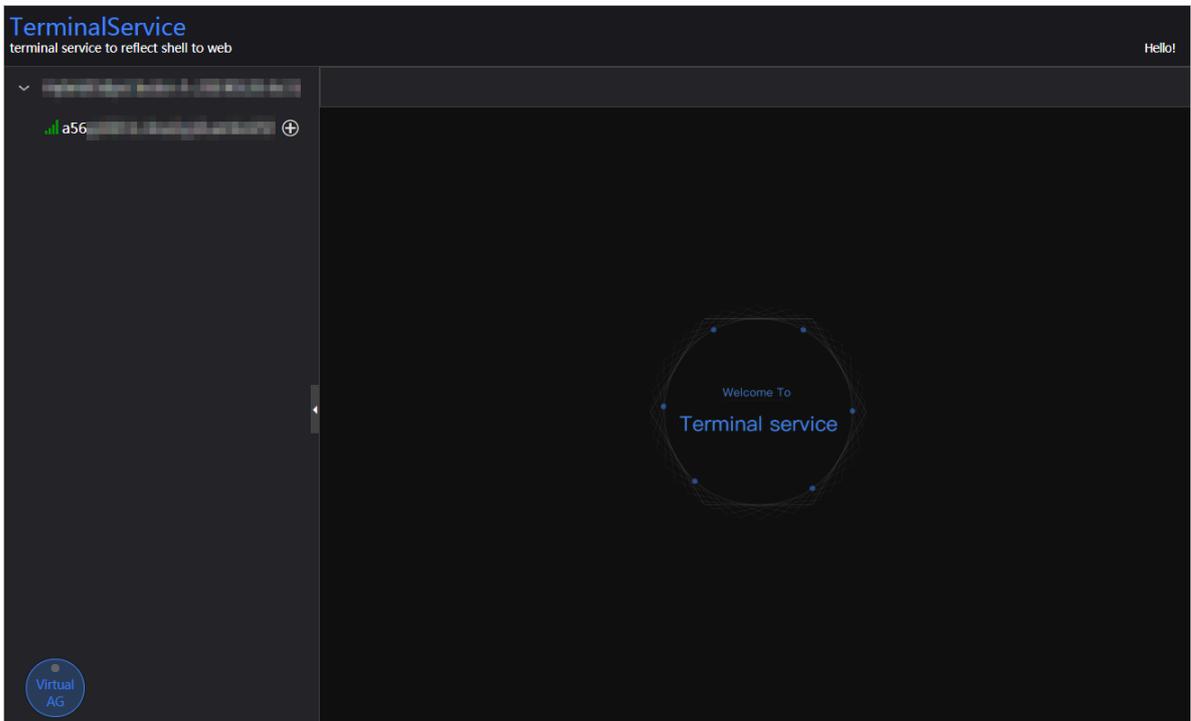
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

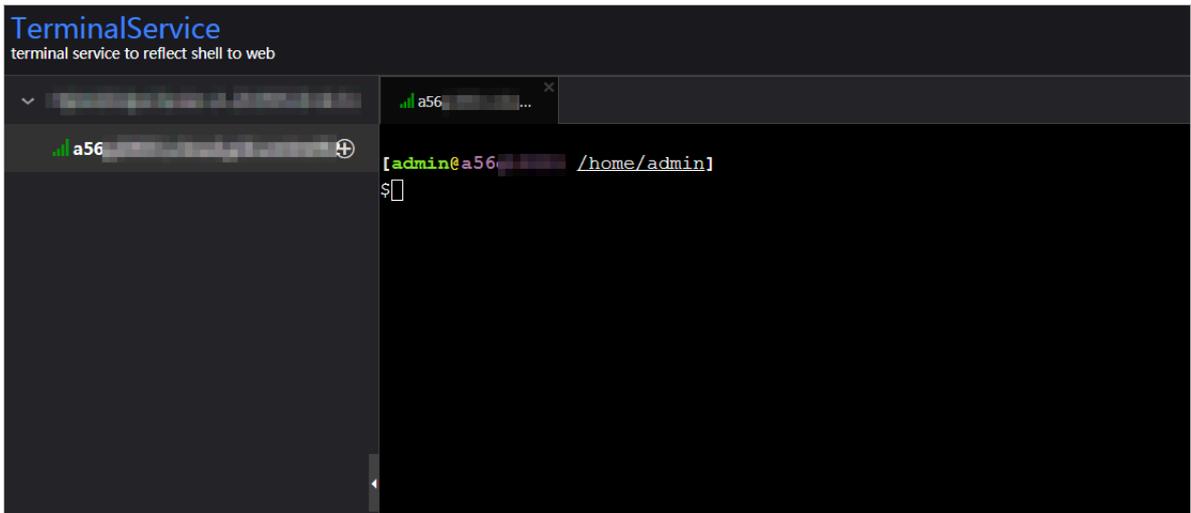
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

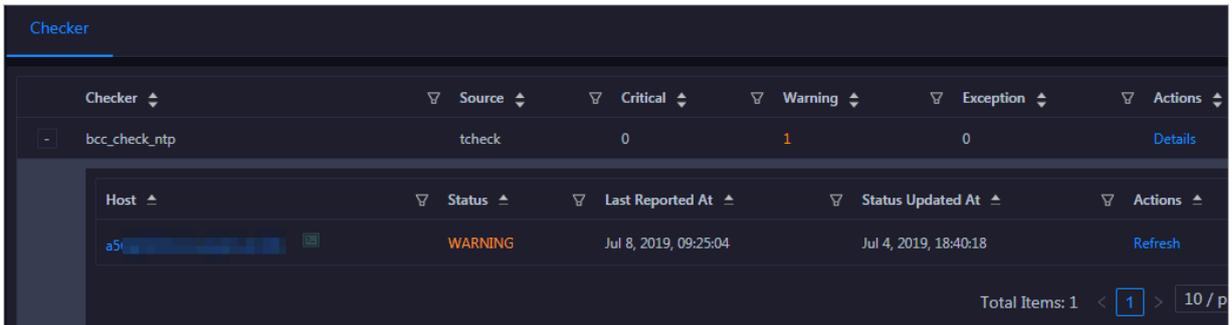


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.10 I+

3.1.10.1 O&M overview

This topic describes the features of I+ O&M and how to access the I+ O&M page.

Modules

I+ O&M includes service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Service O&M	Service overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission , TCP connection, and root disk usage for each service in a cluster.
	Service hosts	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.
Cluster O&M	Cluster overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Cluster health	Displays the check results for a cluster. The check results are divided into the Critical, Warning, Exception, and OK types.
Host O&M	Host overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Host health	Displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click I+.

3. On the page that appears, click O&M at the top. The Services page appears.



The O&M page includes three modules, namely, Services, Clusters, and Hosts.

3.1.10.2 Service O&M

3.1.10.2.1 Service overview

The service overview page lists all I+ services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.

Entry

On the Services page, search for a cluster in the search box above the left-side service list, select a service in the service list, and then click the Overview tab. The Overview page for the service appears.



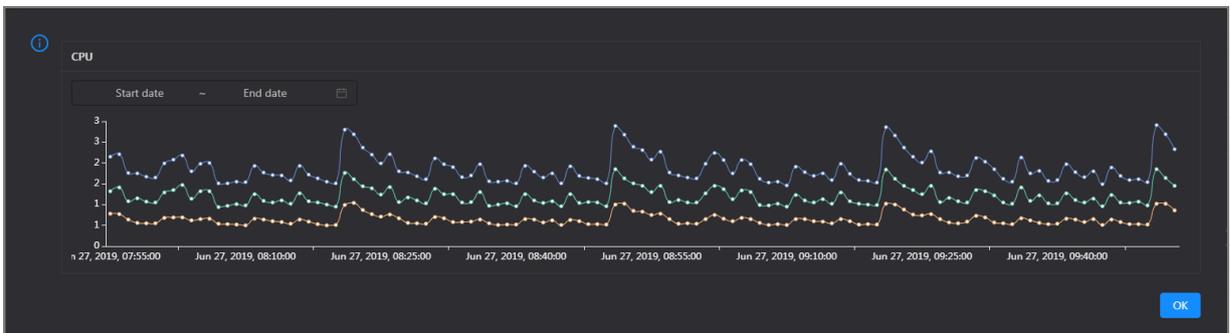
On the Overview page, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

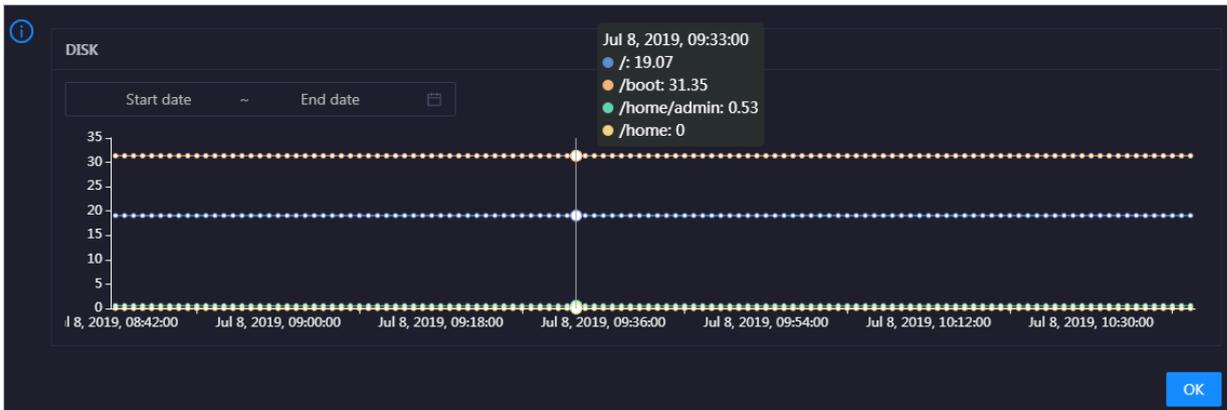
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the service in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the /, /boot, /home/admin, and /home directories for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

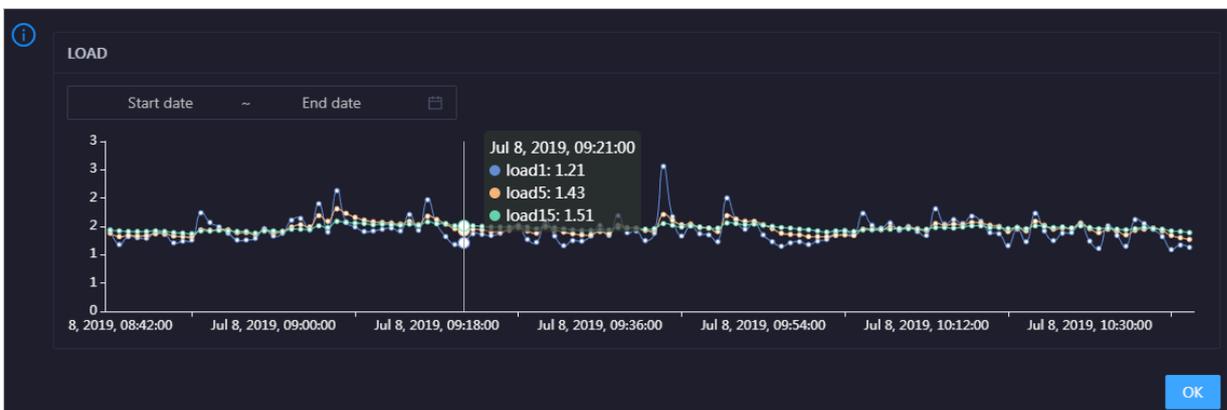


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

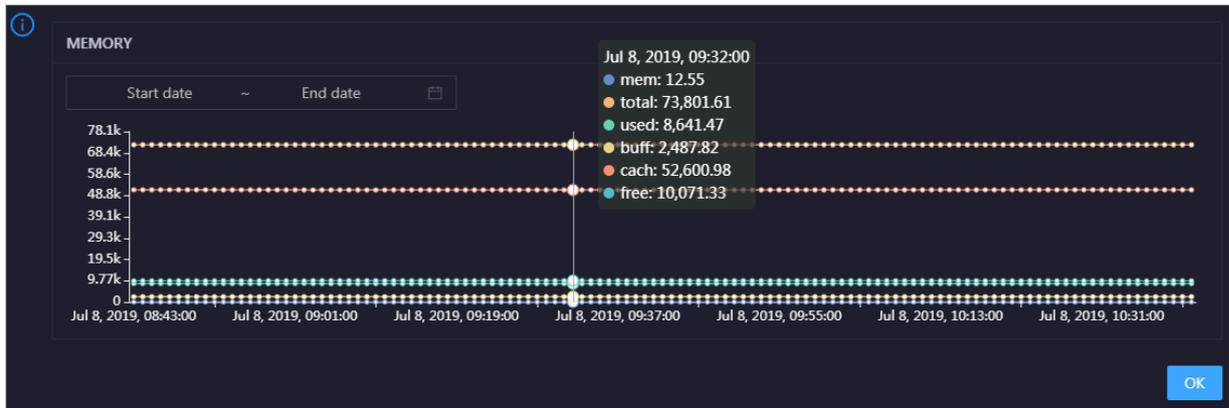


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

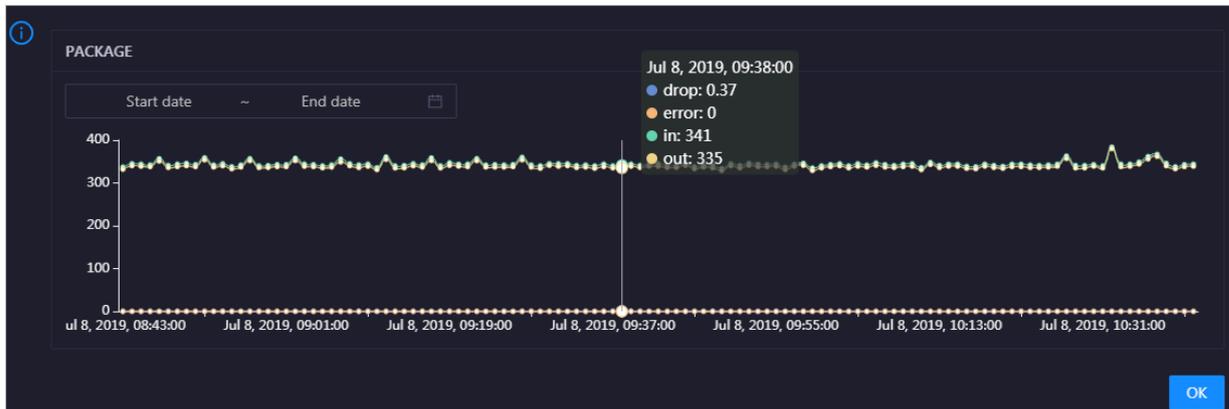


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the selected service over time in different colors. These trend lines reflect the data transmission status of the service.

Click  in the upper-right corner of the chart to zoom in it.



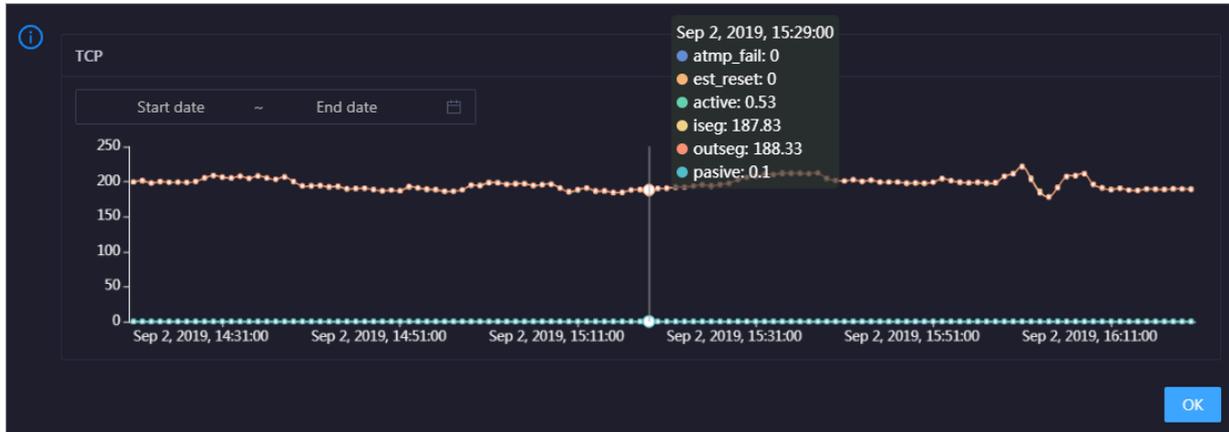
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP

connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the selected service over time in different colors. These trend lines reflect the TCP connection status of the service.

Click  in the upper-right corner of the chart to zoom in it.

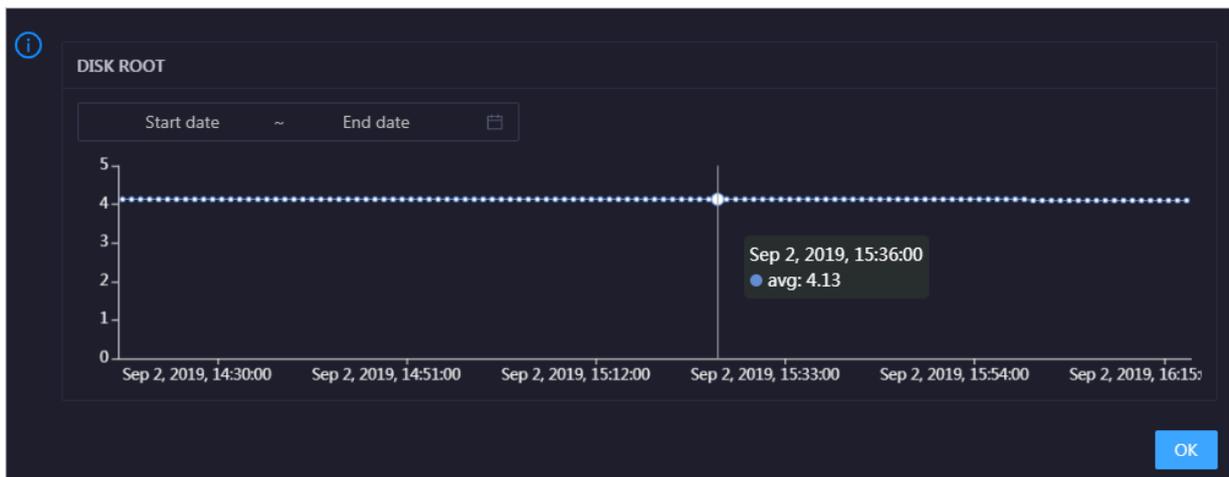


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

This chart displays the trend line of the average root disk usage (avg) for the selected service over time.

Click  in the upper-right corner of the chart to zoom in it.

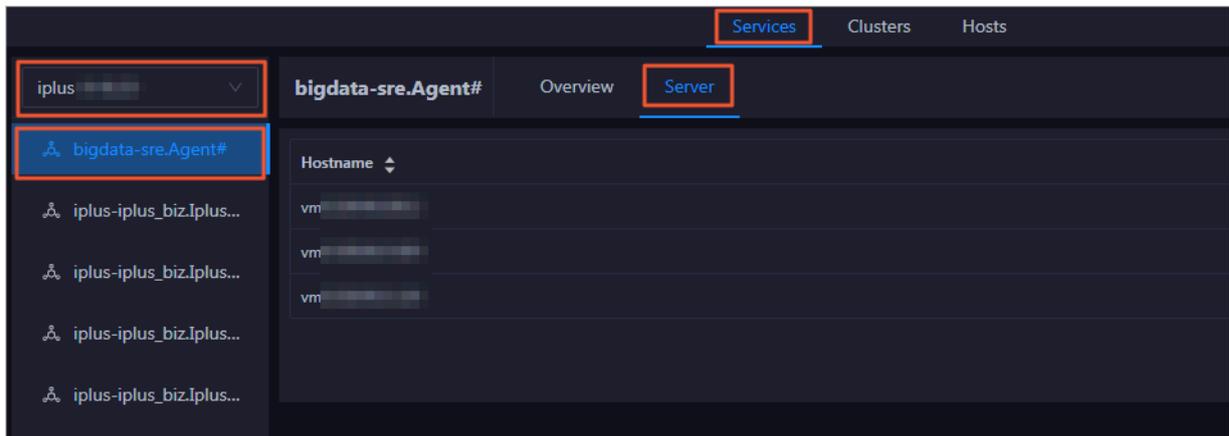


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

3.1.10.2.2 Service hosts

Apsara Bigdata Manager (ABM) allows you to view the host list of each I+ service so that you can understand the service deployment on hosts.

On the Services page, search for a cluster in the search box above the left-side service list, select a service in the service list, and then click the Server tab. The Server page for the service appears.



On the Server page, you can view the hosts where the selected service is run.

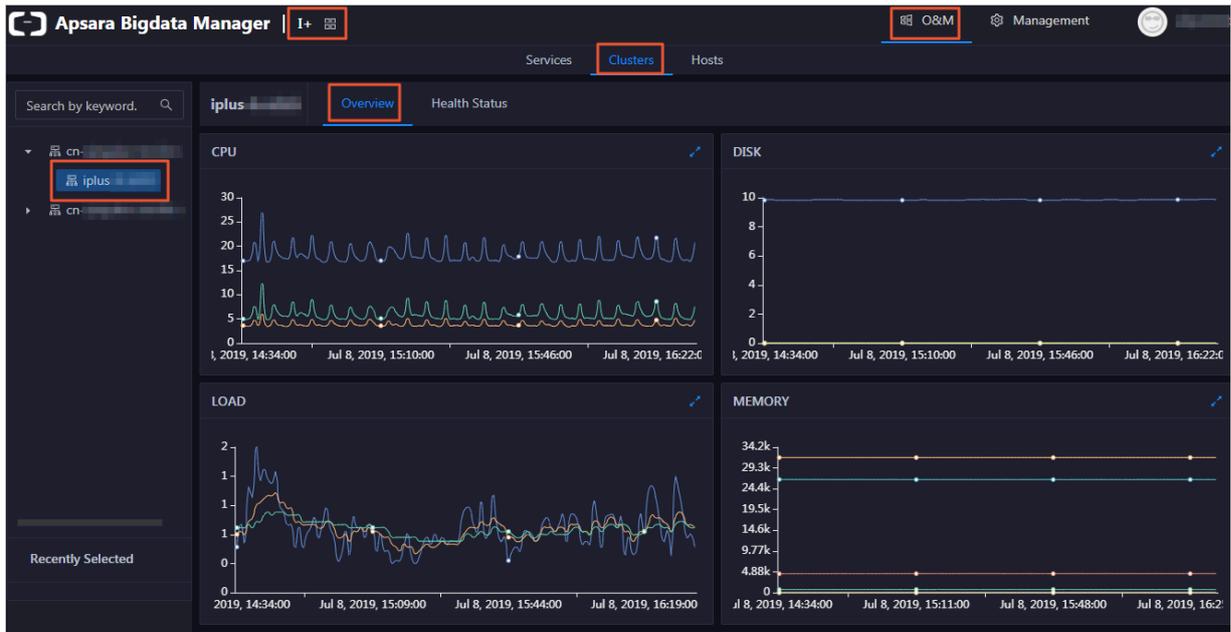
3.1.10.3 Cluster O&M

3.1.10.3.1 Cluster overview

The cluster overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.

Entry

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.



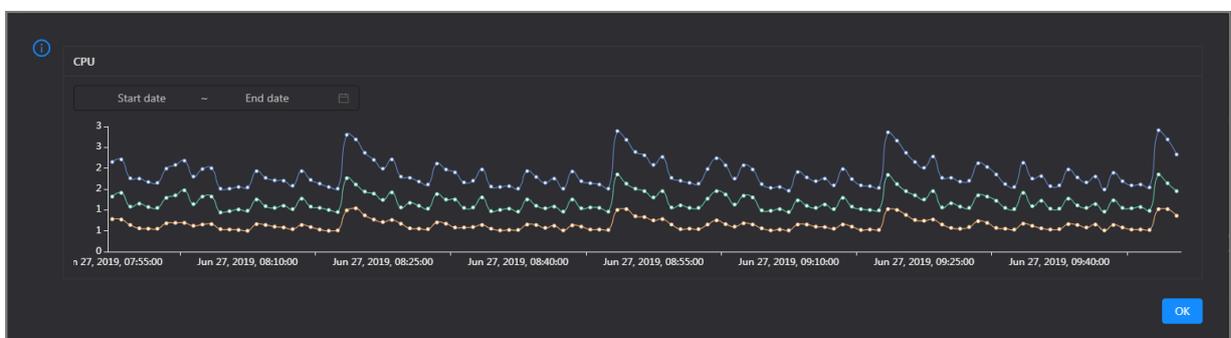
The Overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the cluster. To view information about a cluster, select a region in the left-side navigation pane, and then select a cluster in the region.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

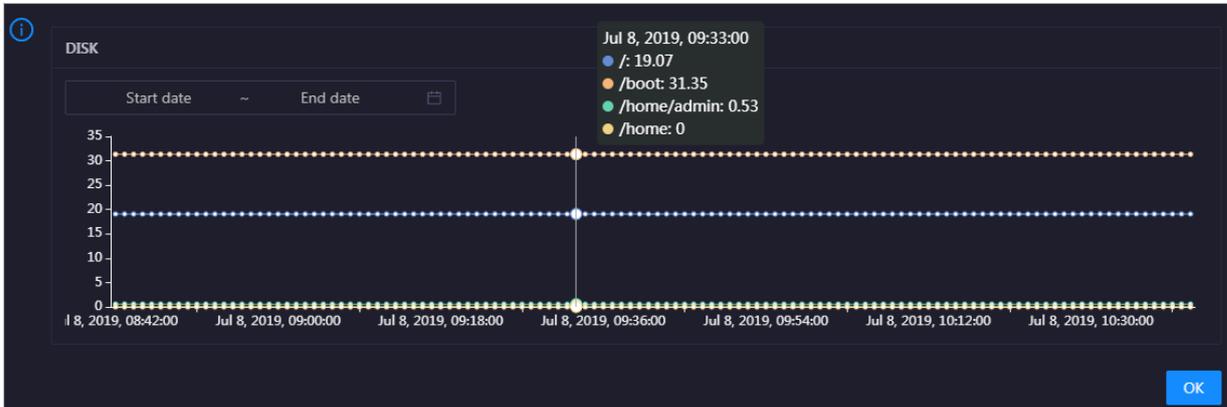
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

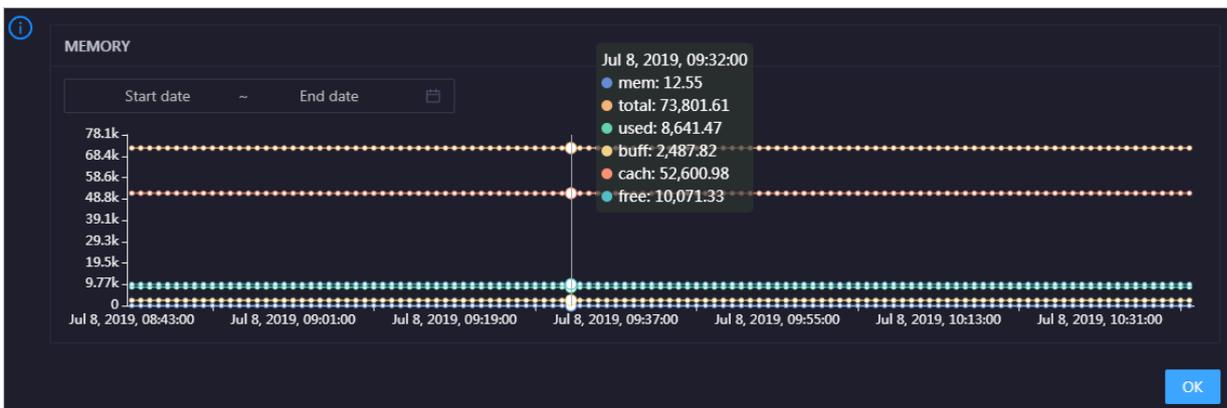


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

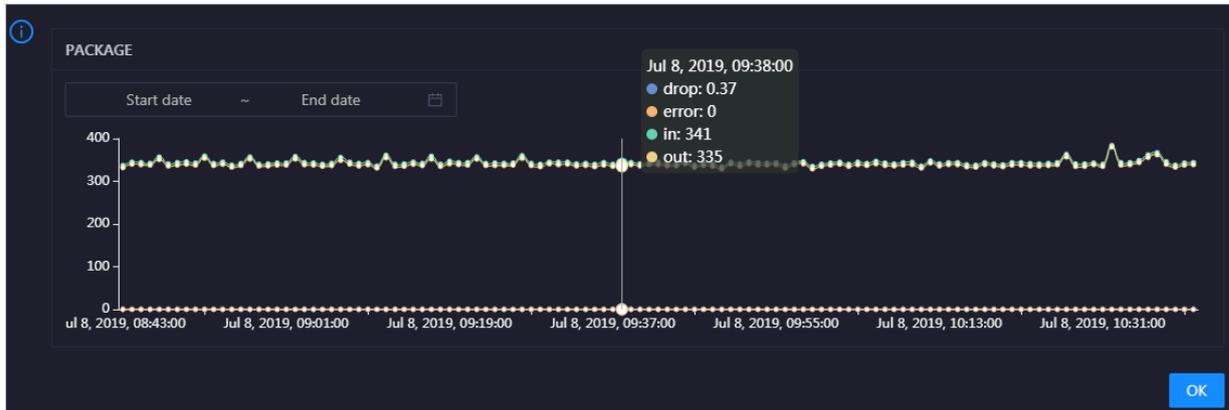


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

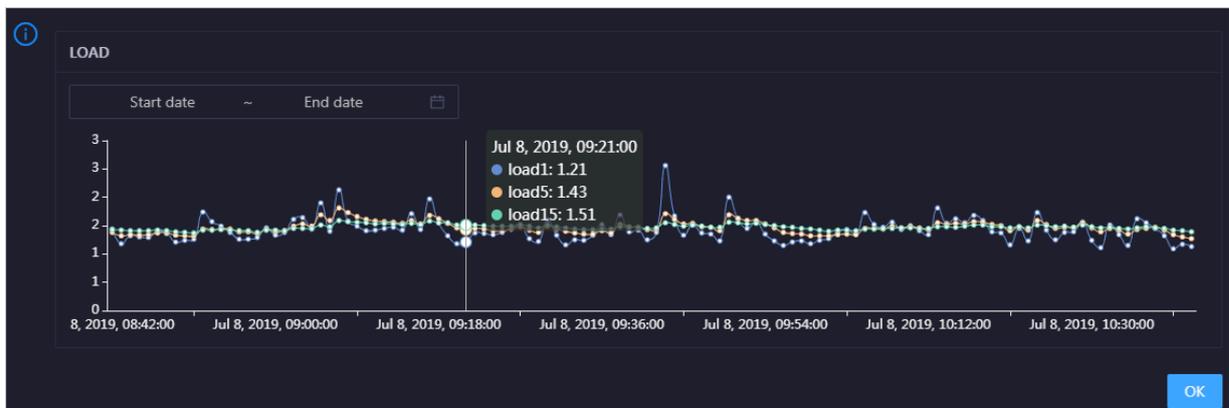


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

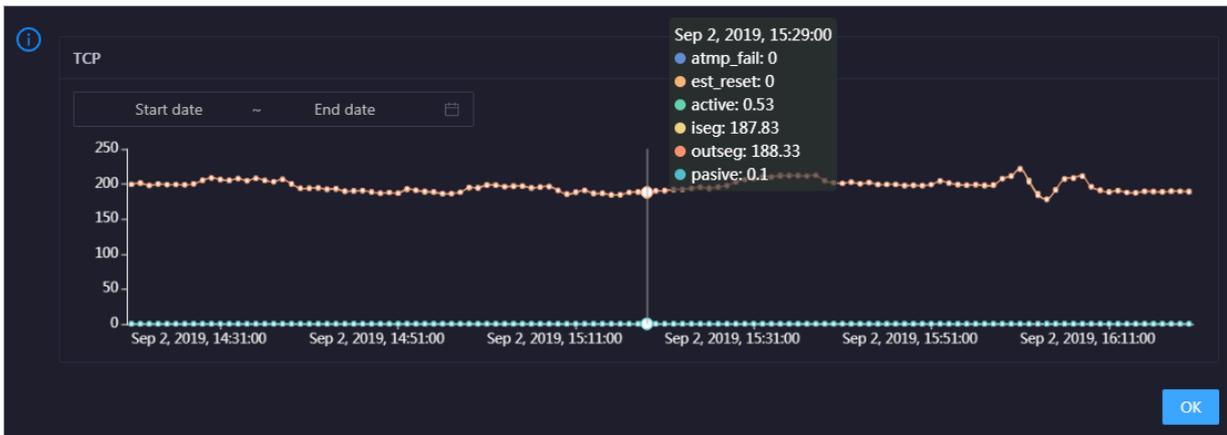


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

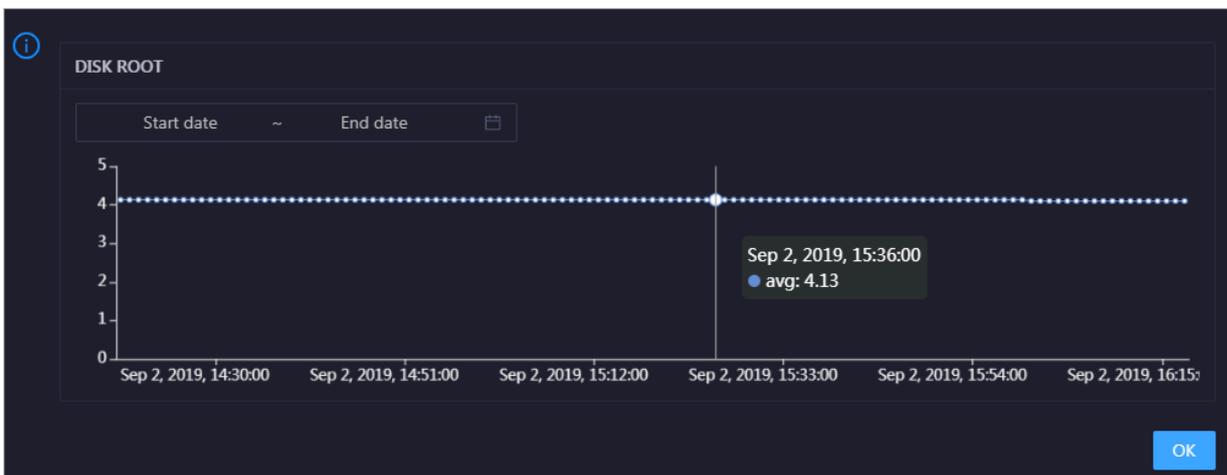


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

Click  in the upper-right corner of the chart to zoom in the chart.



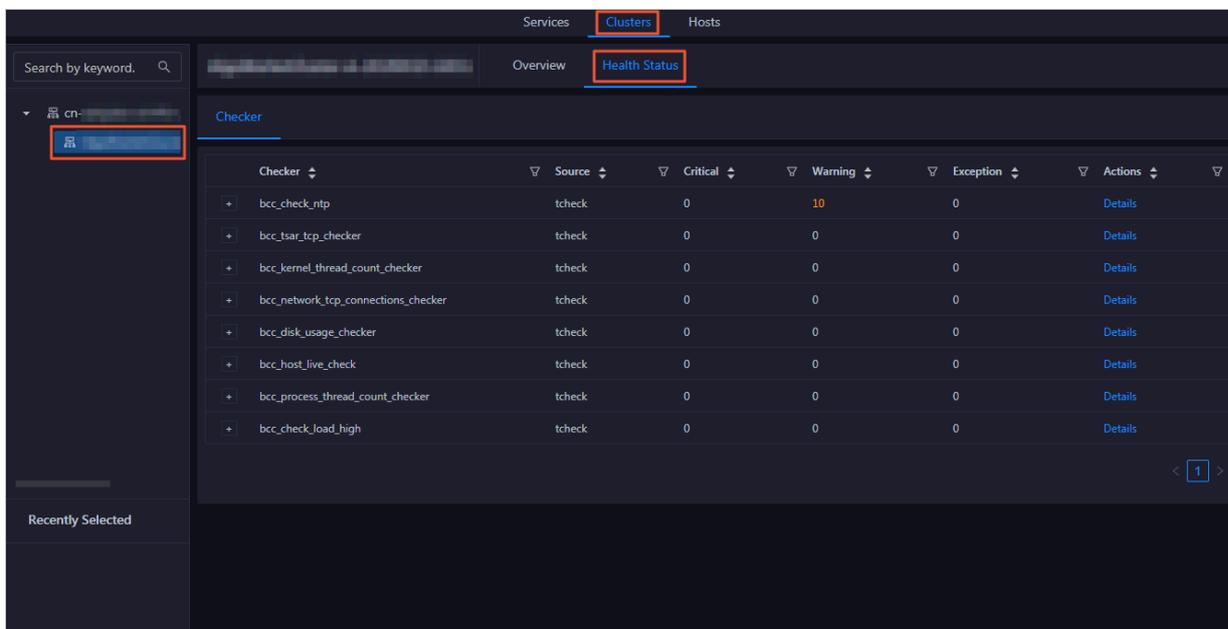
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

3.1.10.3.2 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab.



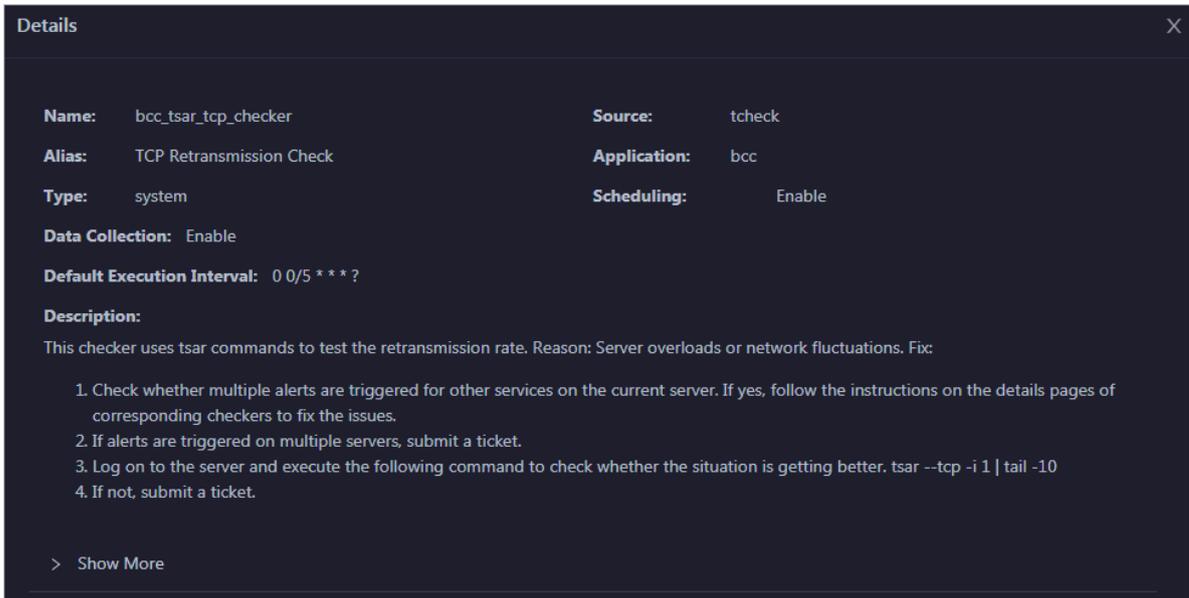
The screenshot shows the 'Clusters' page with the 'Health Status' tab selected. A table lists various checkers with their status. The 'Warning' column for 'bcc_check_ntp' is highlighted in orange with the value '10'. The 'Critical' and 'Exception' columns for all checkers show '0'. The 'Actions' column contains a 'Details' link for each checker.

Checker	Source	Critical	Warning	Exception	Actions
+ bcc_check_ntp	tcheck	0	10	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_network_tcp_connections_checker	tcheck	0	0	0	Details
+ bcc_disk_usage_checker	tcheck	0	0	0	Details
+ bcc_host_live_check	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

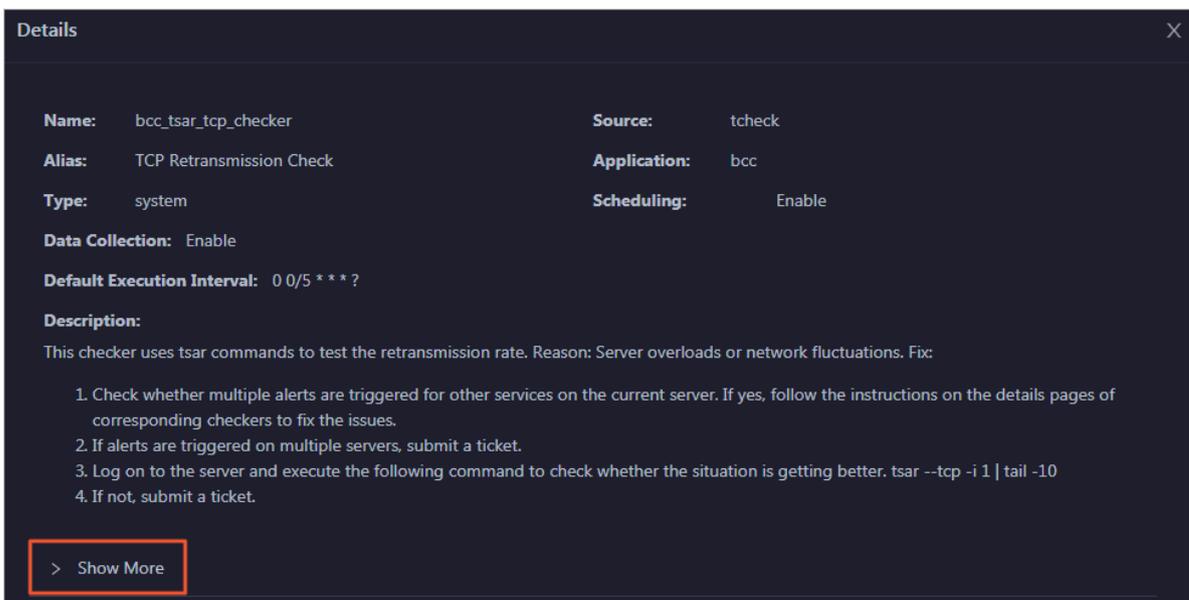
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

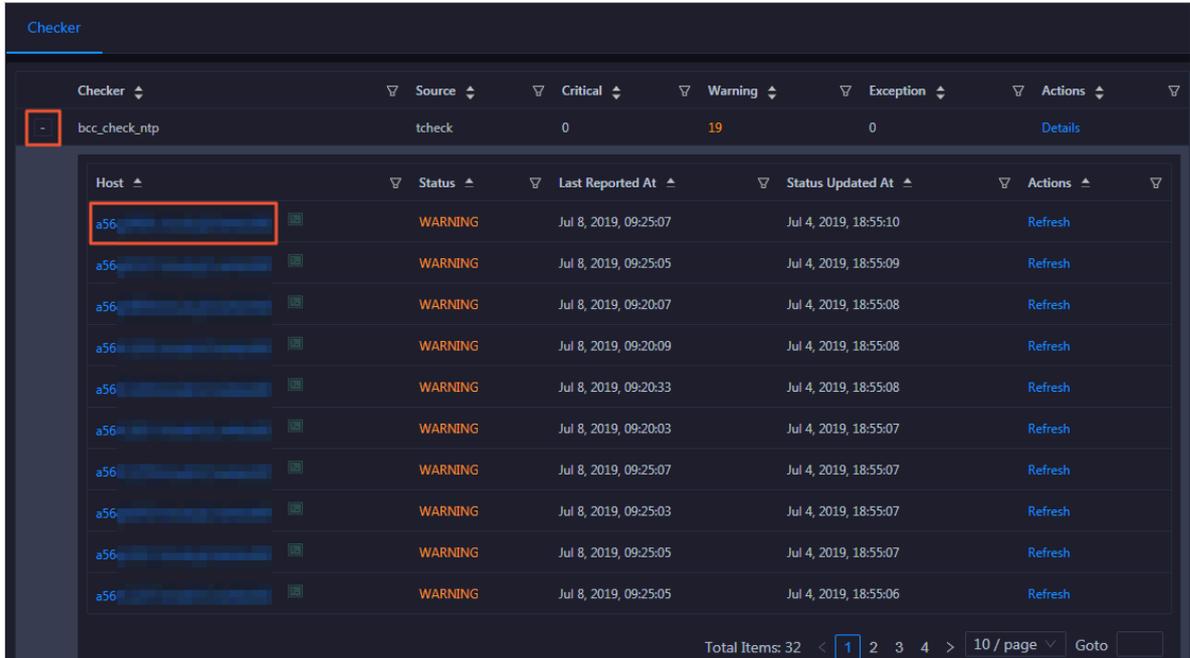


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

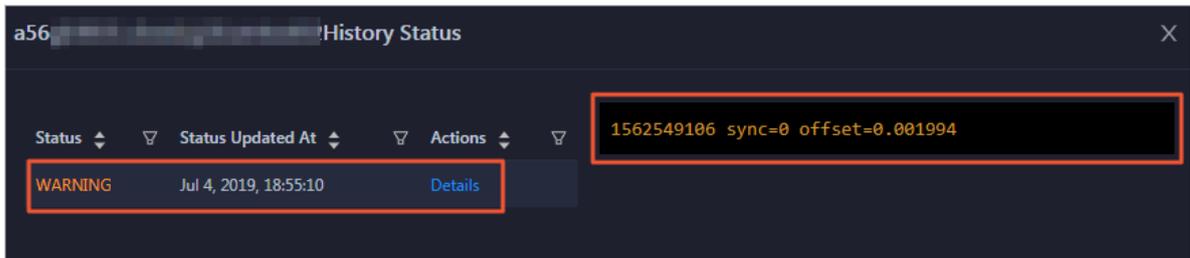
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

- 1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.**

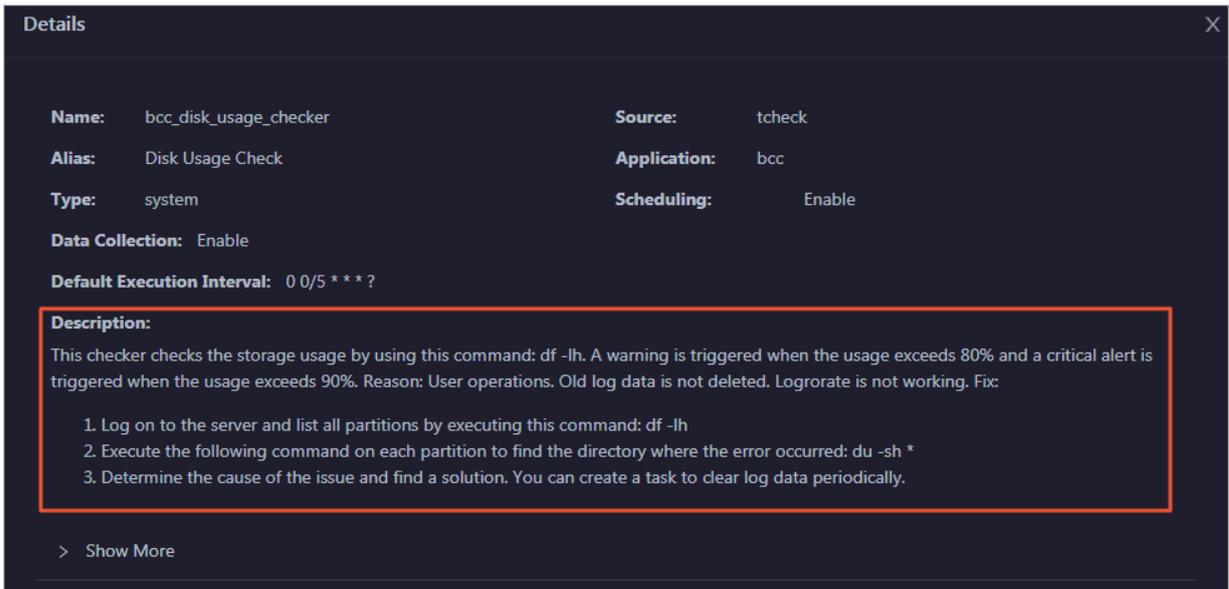


- 2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.**



Clear alerts

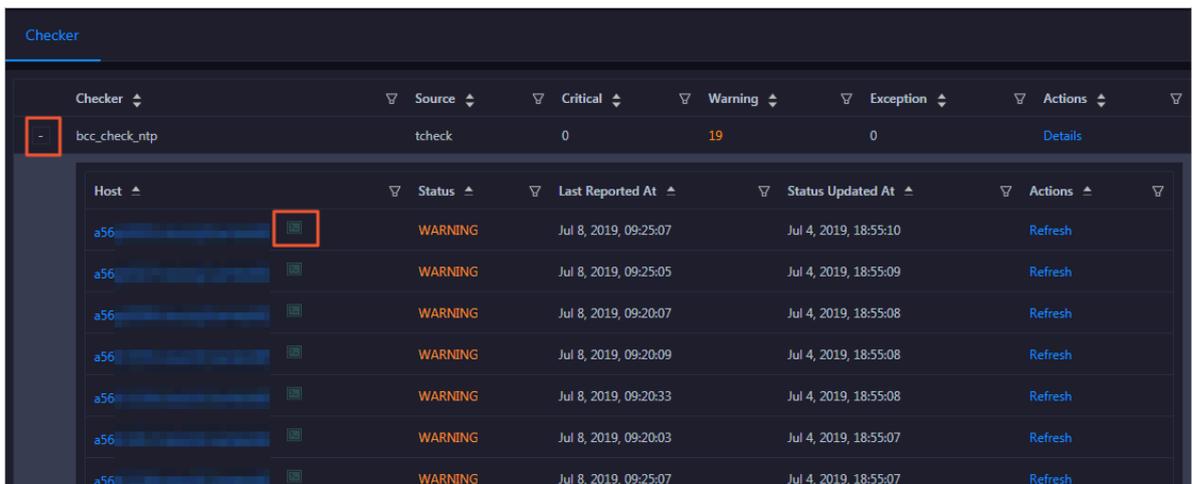
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



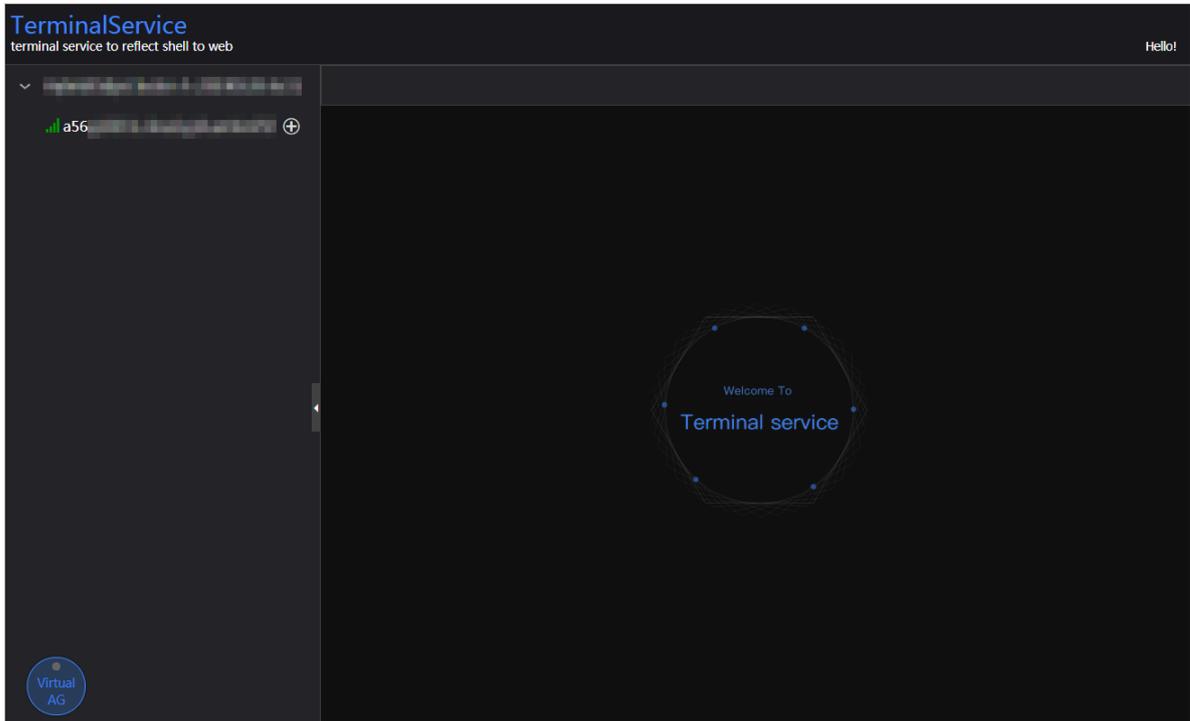
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

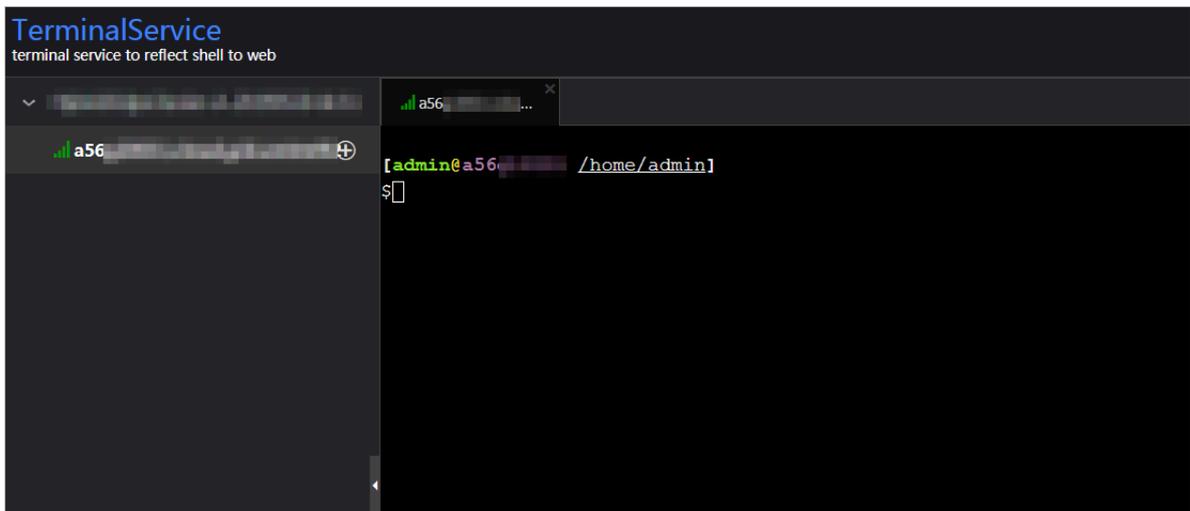
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.

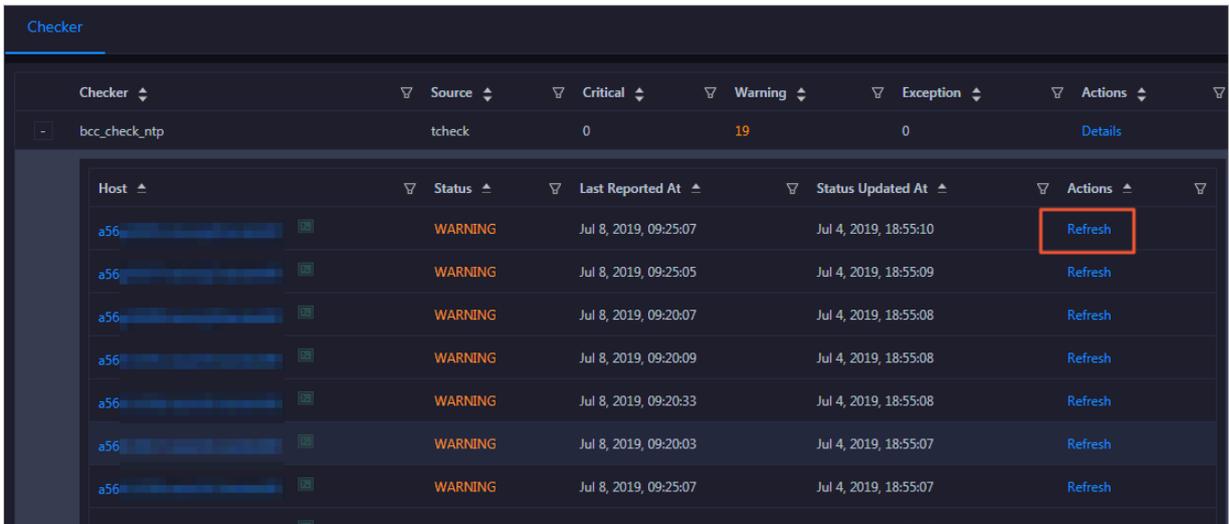


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



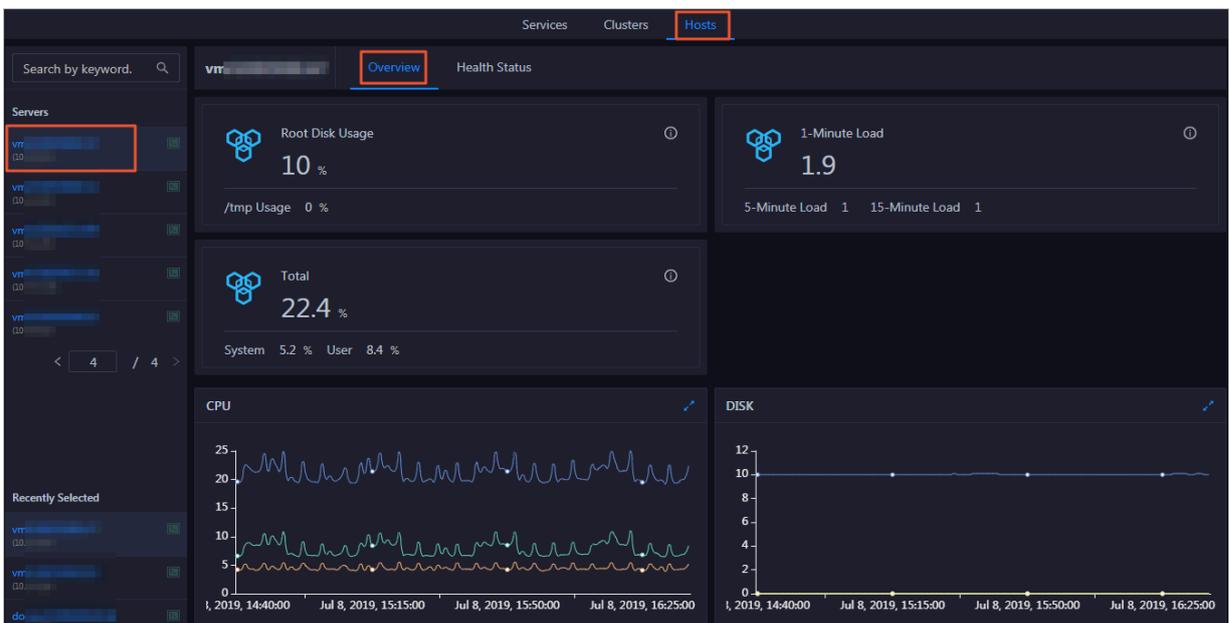
3.1.10.4 Host O&M

3.1.10.4.1 Host overview

The host overview page displays the overall running information about a host in an I+ cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Entry

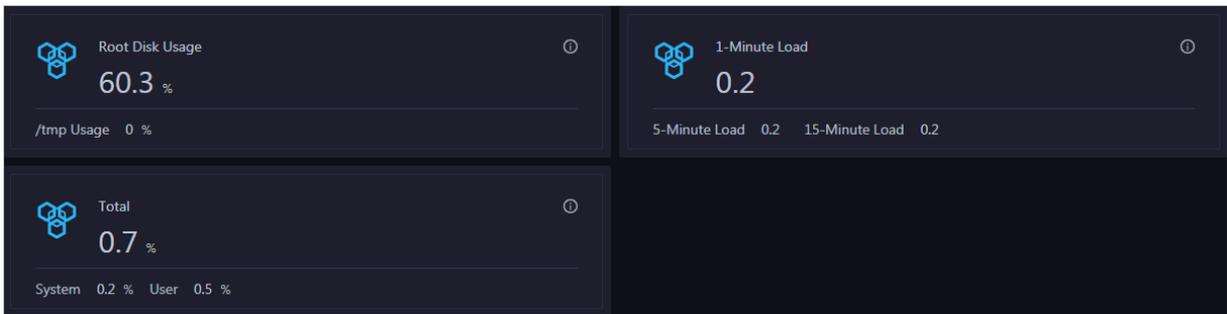
On the Hosts page, select a host in the left-side navigation pane. The Overview page for the host appears.



On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Root Disk Usage, Total, and 1-Minute Load

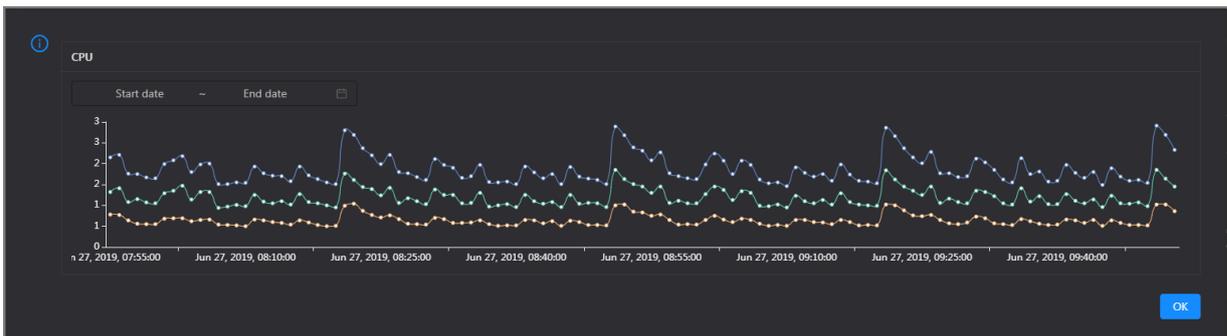
These sections display the root disk usage, total usage, and 1-minute load for the selected host. The Root Disk Usage section provides the usage of the `/tmp` directory. The Total section provides the system usage and user usage. The 1-Minute Load section provides the 1-minute, 5-minute, and 15-minute load averages.



CPU

This chart displays the trend lines of the total CPU usage (`cpu`), CPU usage for executing code in kernel space (`sys`), and CPU usage for executing code in user space (`user`) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

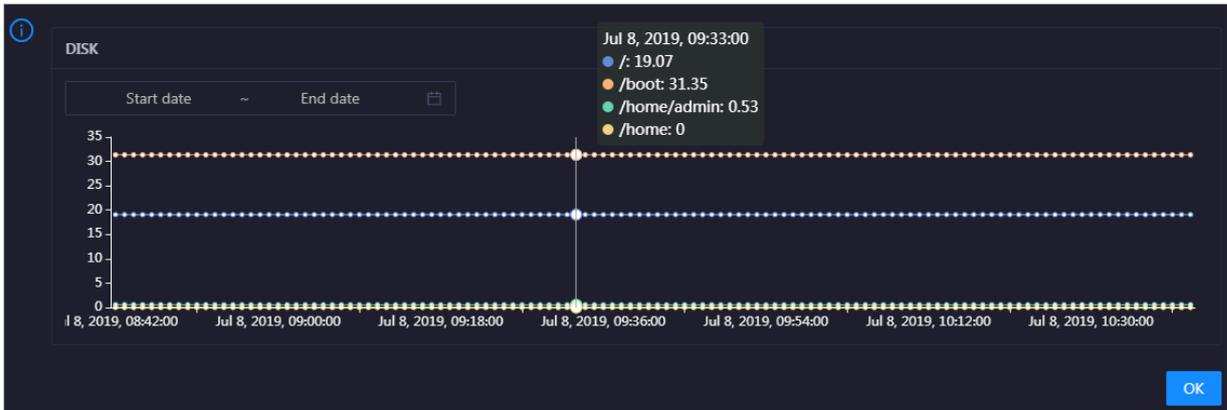


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

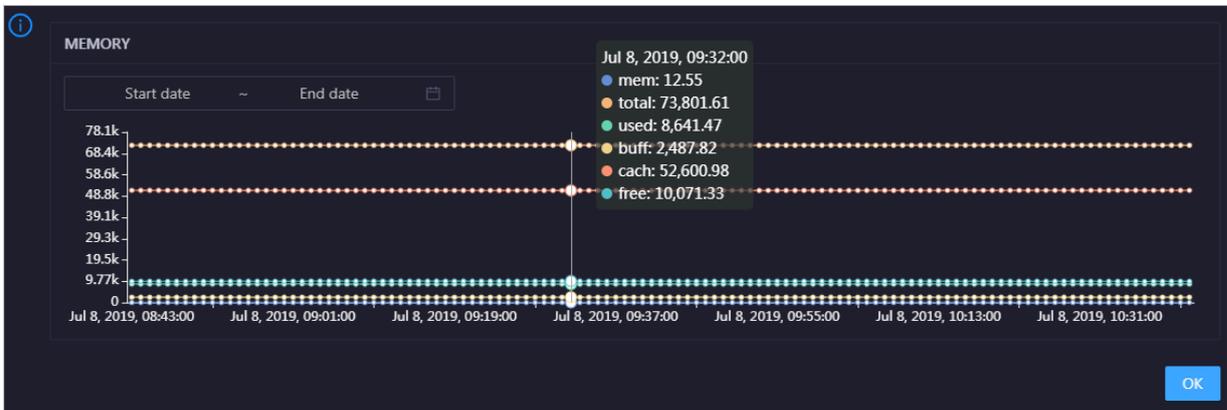


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

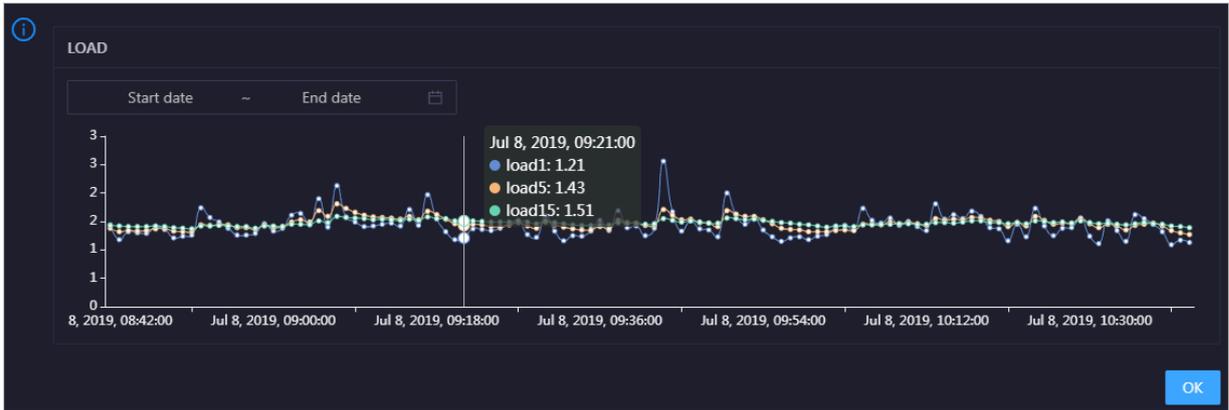


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

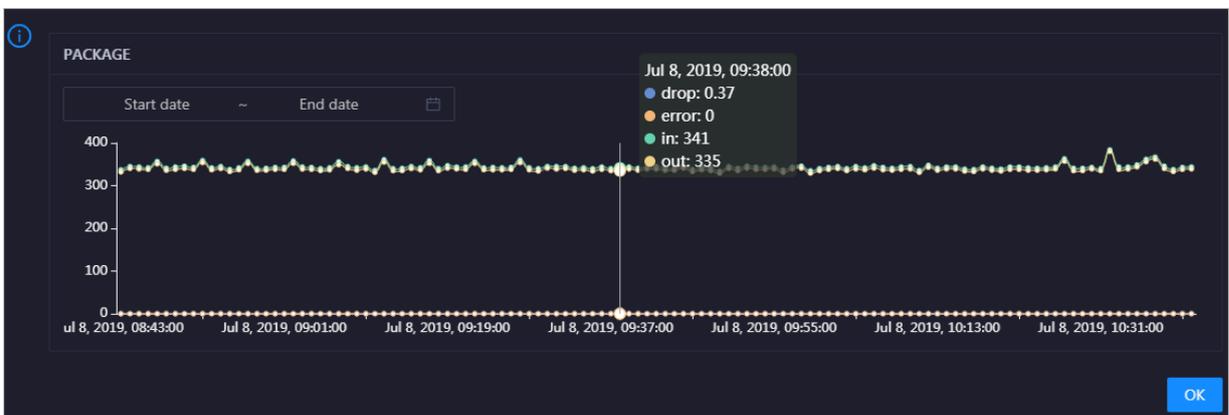


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.



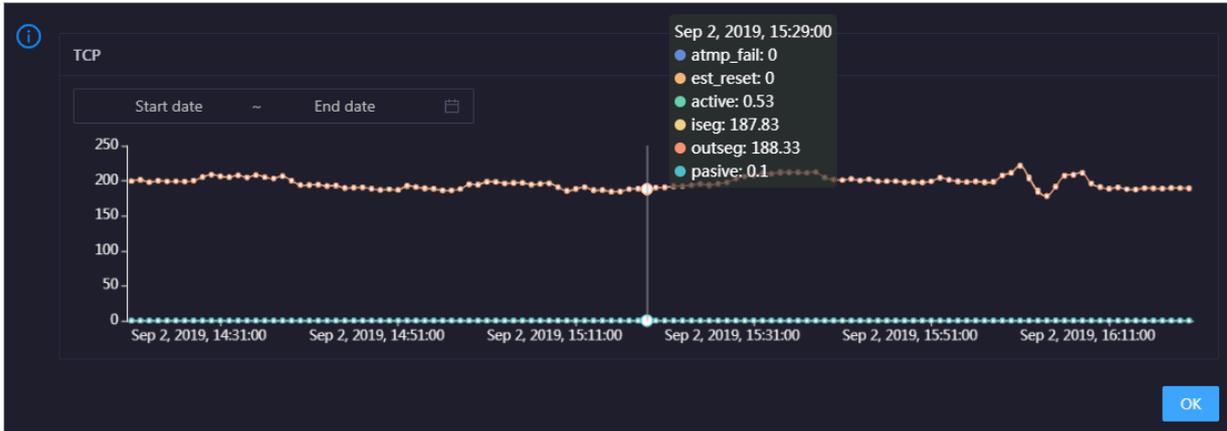
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of

sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click  in the upper-right corner of the chart to zoom in the chart.

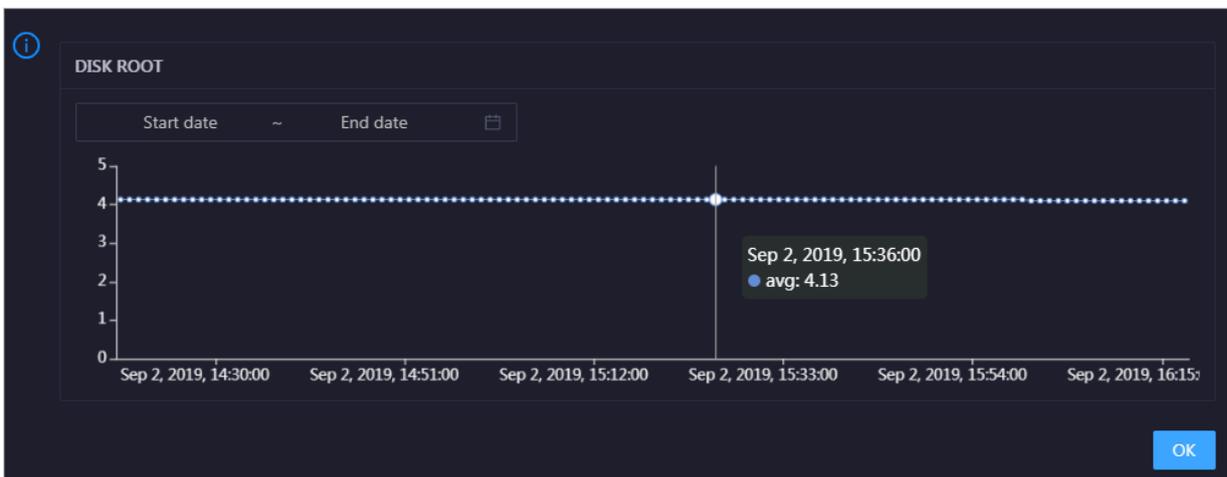


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the host over time.

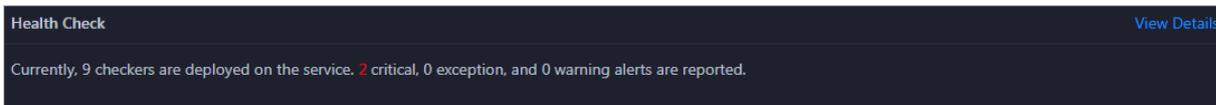
Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

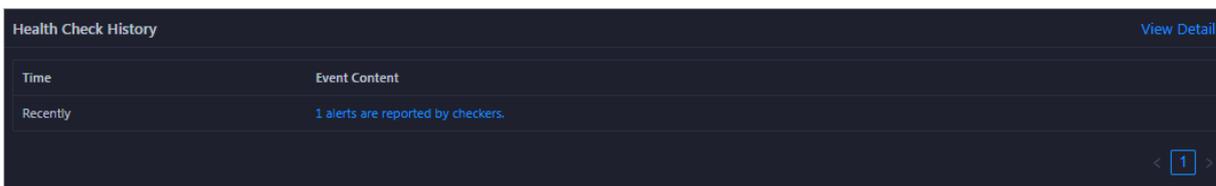
This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.



Click View Details to go to the [Host health](#) page. On this page, you can view the health check details.

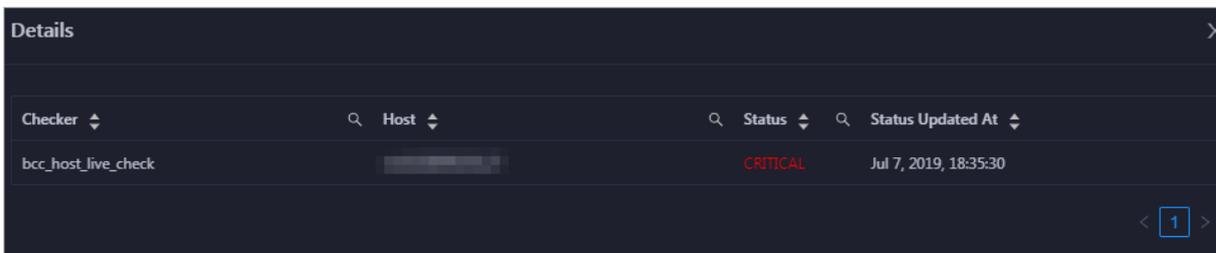
Health Check History

This section displays a record of the health checks performed on the host.



Click View Details to go to the [Host health](#) page. On this page, you can view the health check details.

You can click the event content of a check to view the exception items.

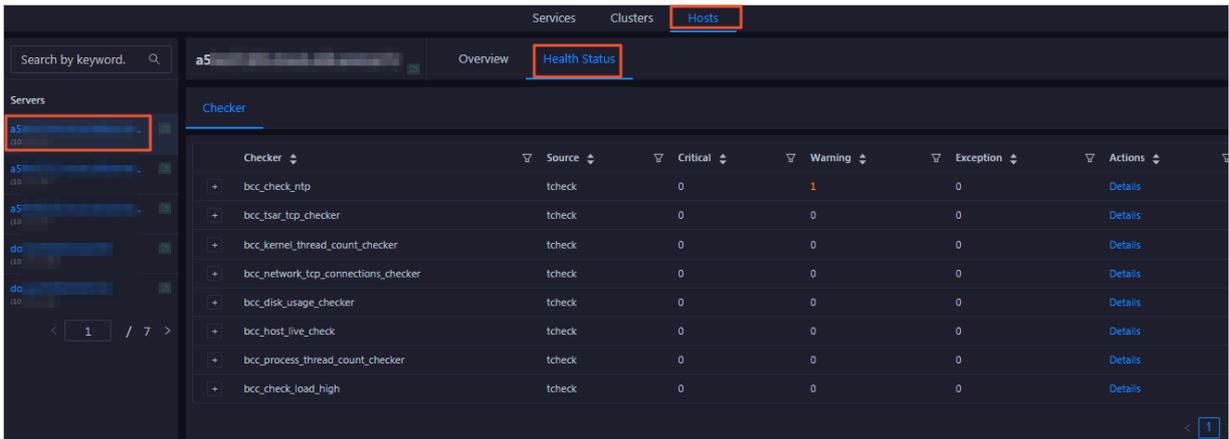


3.1.10.4.2 Host health

On the host health status page, you can view the checkers of all hosts, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

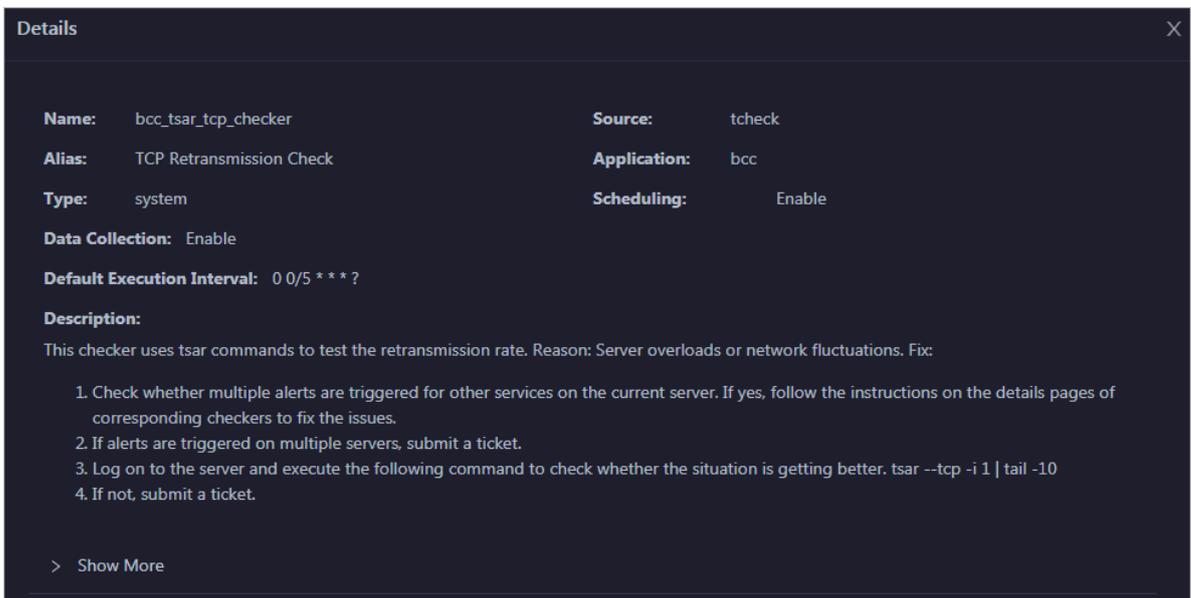
On the Hosts page, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.



On the Health Status page, you can view all checkers of the host and the check results for the hosts in the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

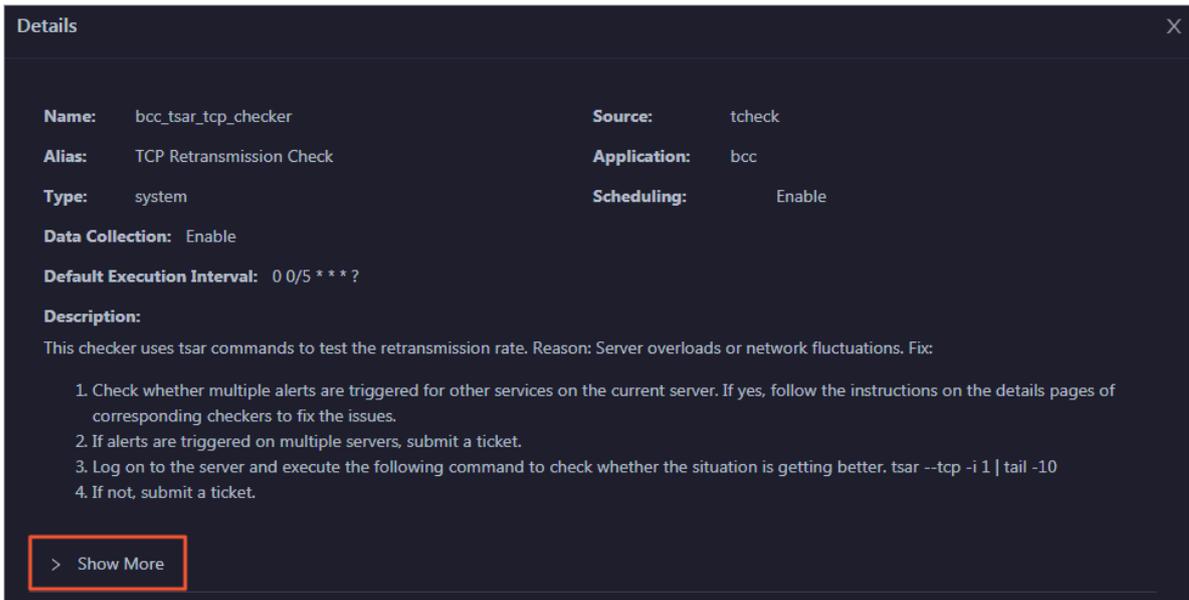
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

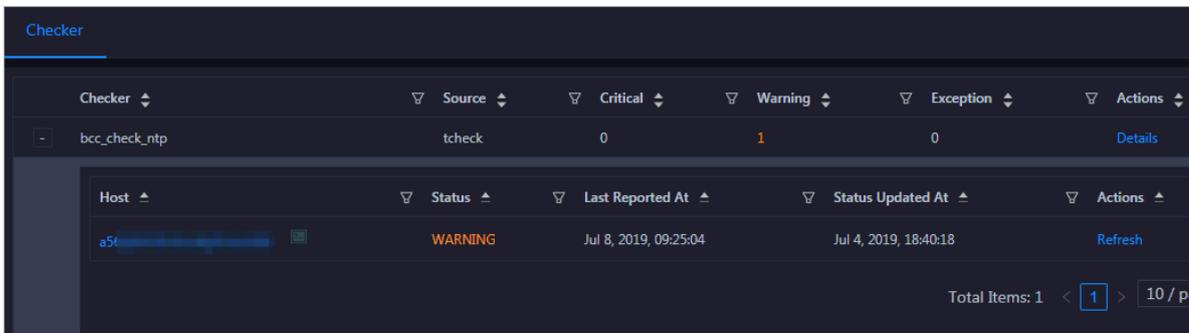


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

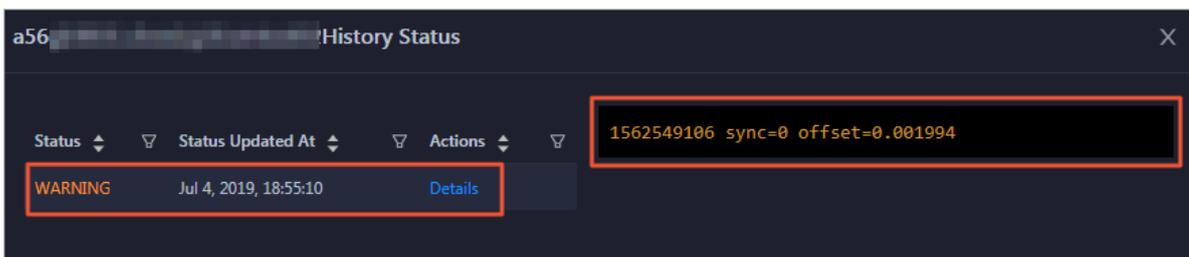
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

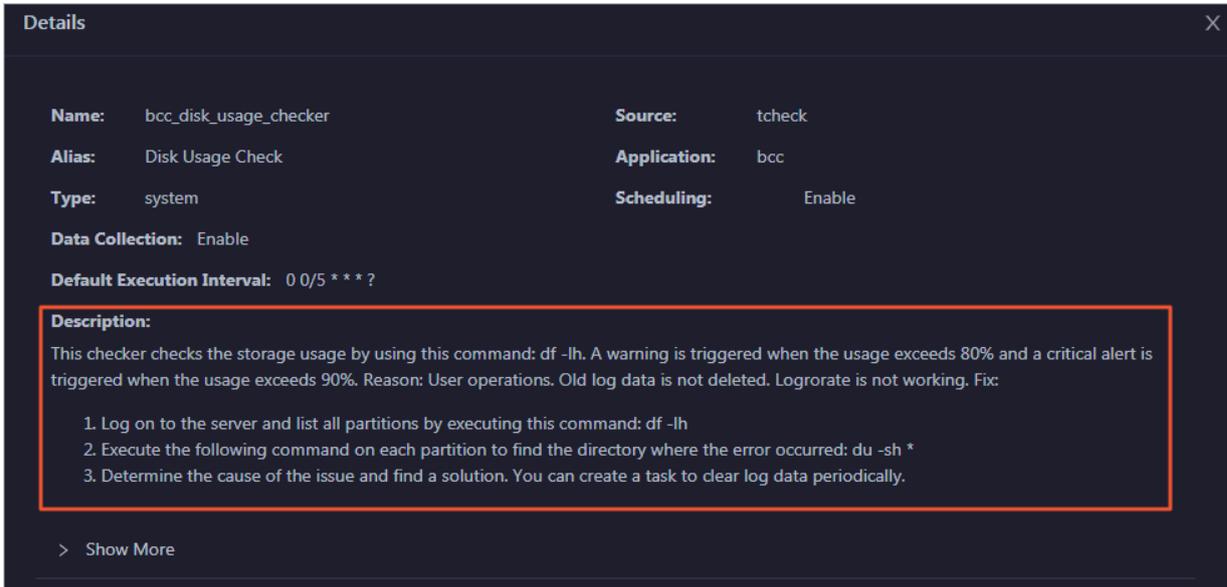


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

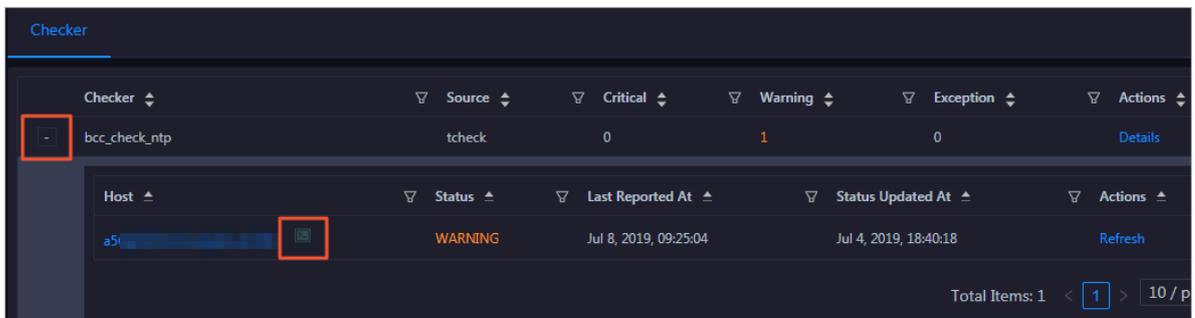
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



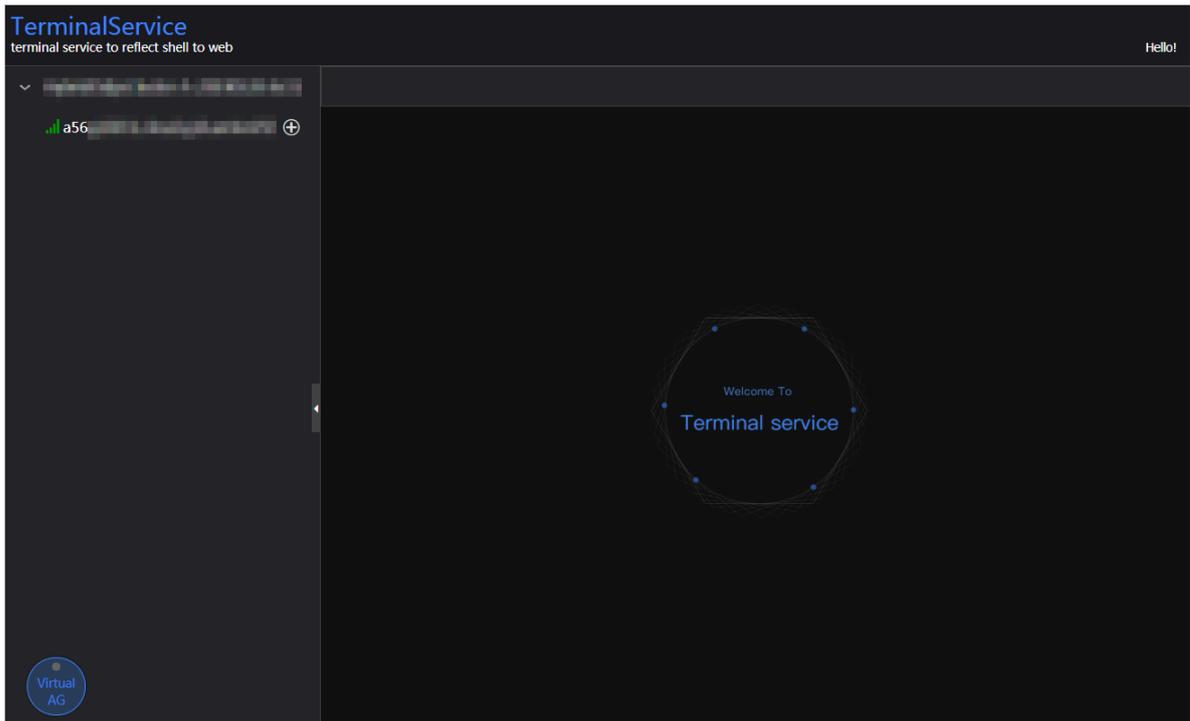
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

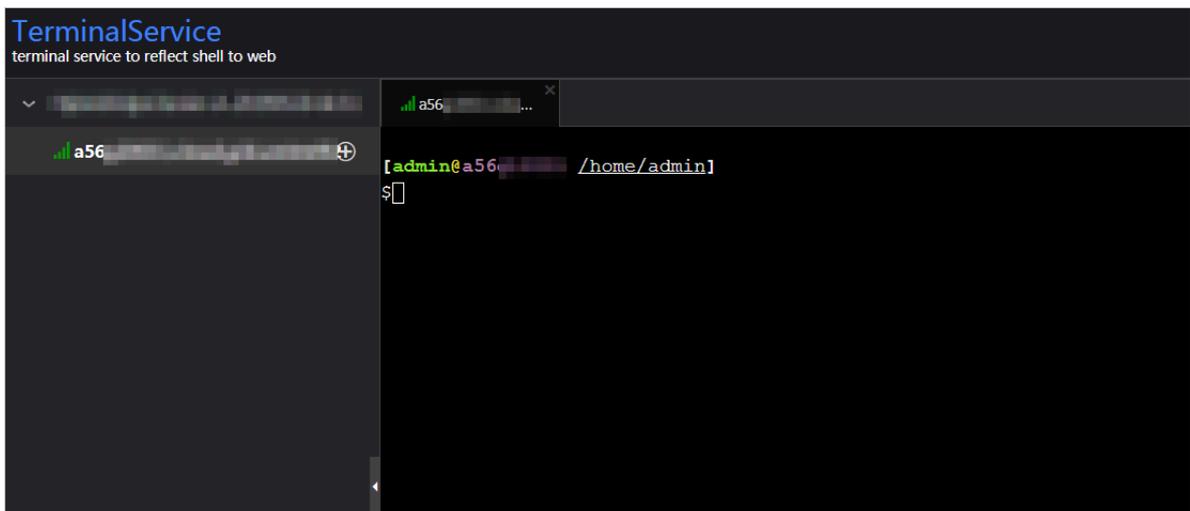
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

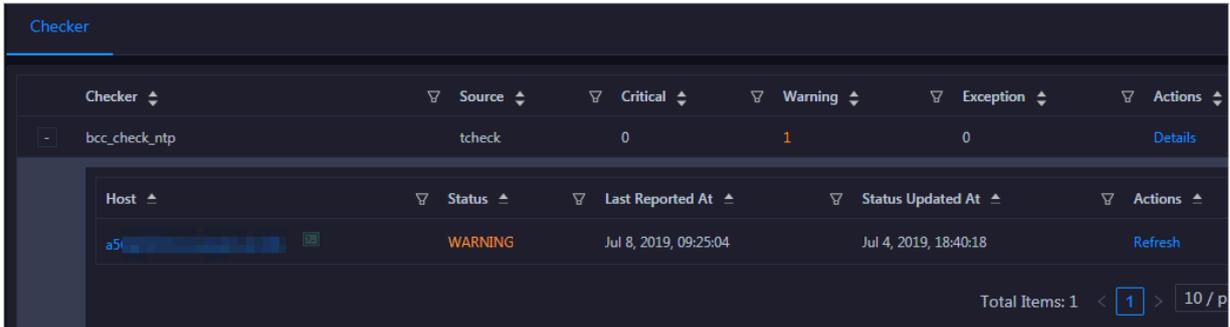


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.11 Quick BI

3.1.11.1 O&M overview

This topic describes the features of Quick BI O&M and how to access the Quick BI O&M page.

Modules

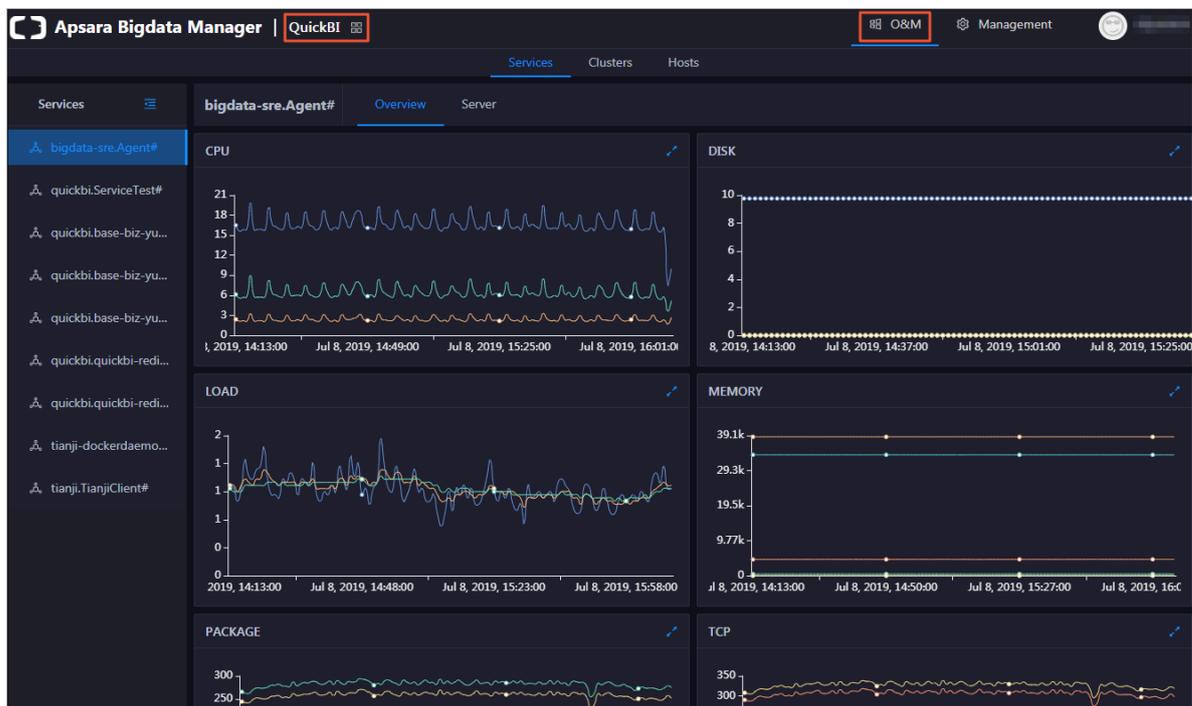
Quick BI O&M includes service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Service O&M	Service overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission , TCP connection, and root disk usage for each service in a cluster.
	Service hosts	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.
Cluster O&M	Cluster overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Cluster health	Displays the check results for a cluster. The check results are divided into the Critical, Warning, Exception, and OK types.

Module	Feature	Description
Host O&M	Host overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Host health	Displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click Quick BI.
3. On the page that appears, click O&M at the top. The Services page appears.



The O&M page includes three modules, namely, Services, Clusters, and Hosts.

3.1.11.2 Service O&M

3.1.11.2.1 Service overview

The service overview page lists all Quick BI services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.

Entry

1. At the top of the O&M page, click Services.
2. On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.
3. Click the Overview tab. The Overview page for the service appears.



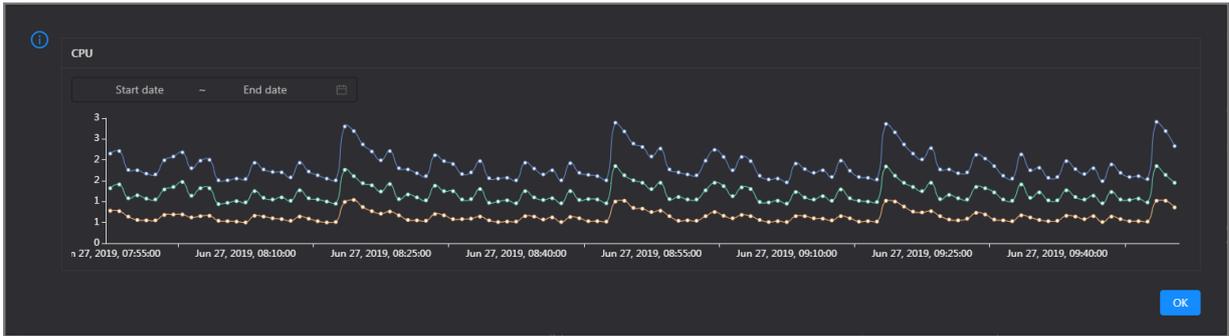
On the Overview page, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

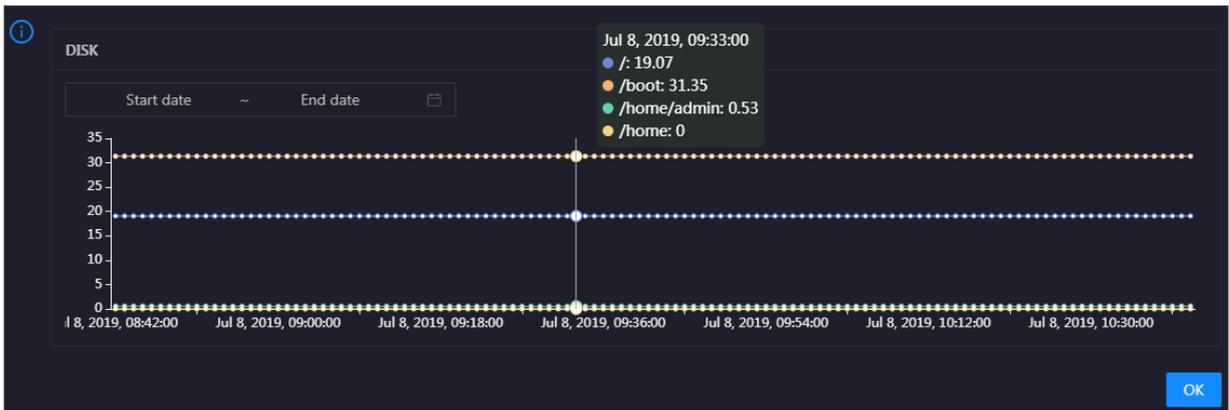
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the service in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

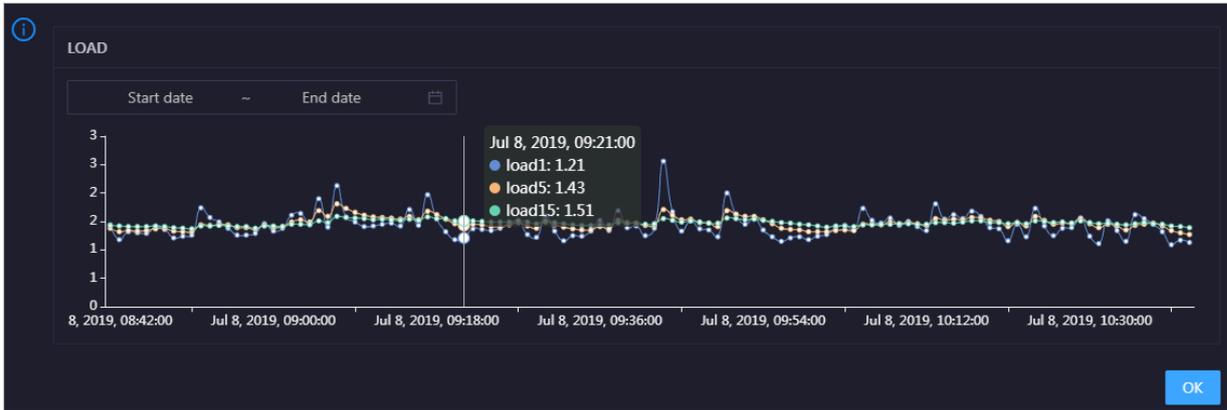


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

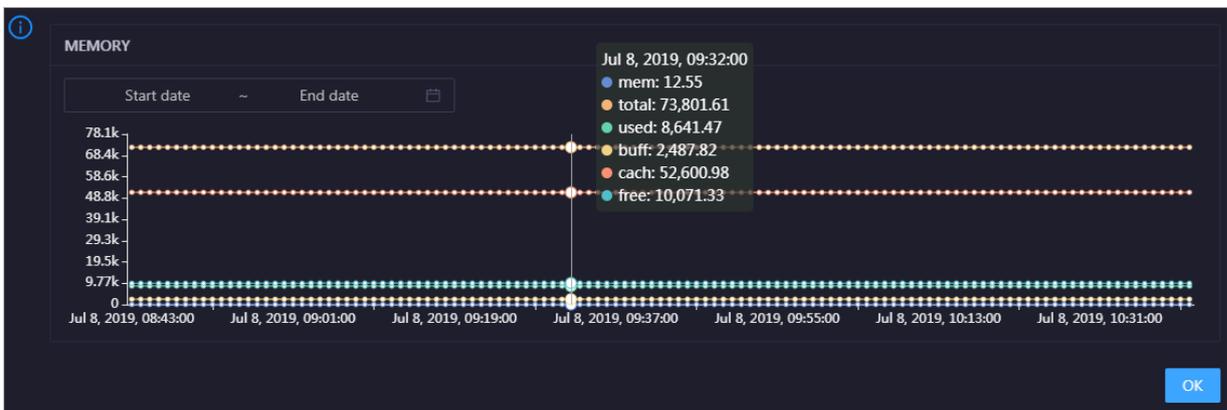


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

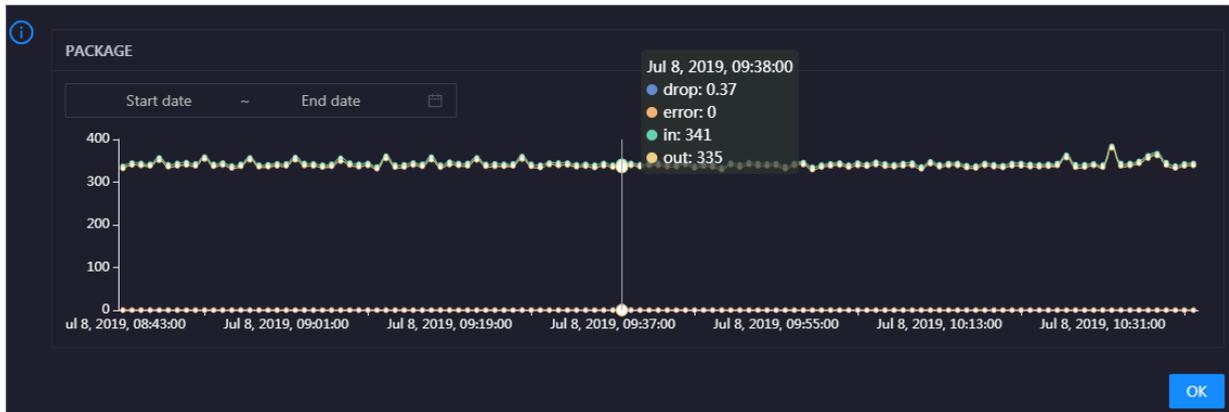


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the selected service over time in different colors. These trend lines reflect the data transmission status of the service.

Click  in the upper-right corner of the chart to zoom in it.

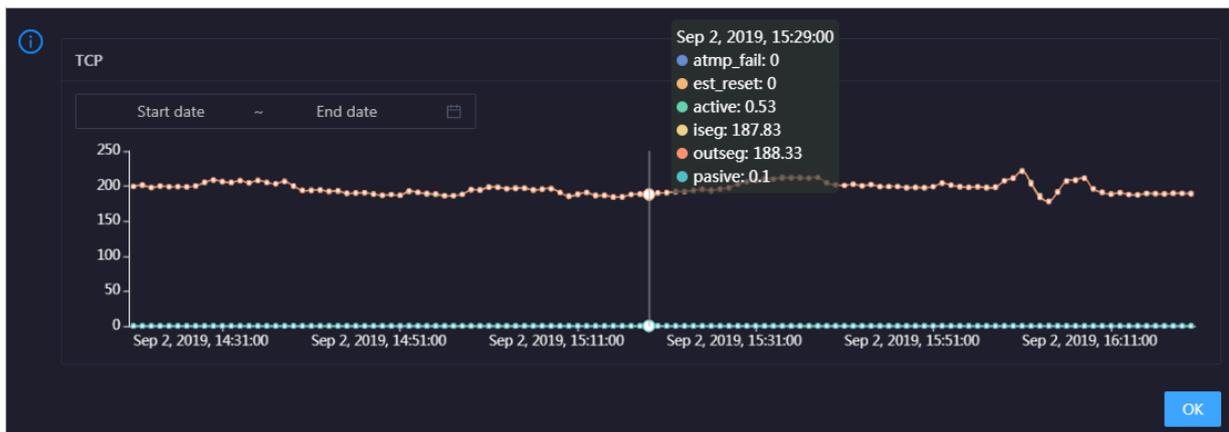


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (`atmp_fail`), that of the times of resetting TCP connections in the ESTABLISHED state (`est_reset`), that of active TCP connections (`active`), that of passive TCP connections (`pasive`), that of received TCP packets (`iseg`), and that of sent TCP packets (`outseg`) for the selected service over time in different colors. These trend lines reflect the TCP connection status of the service.

Click  in the upper-right corner of the chart to zoom in it.

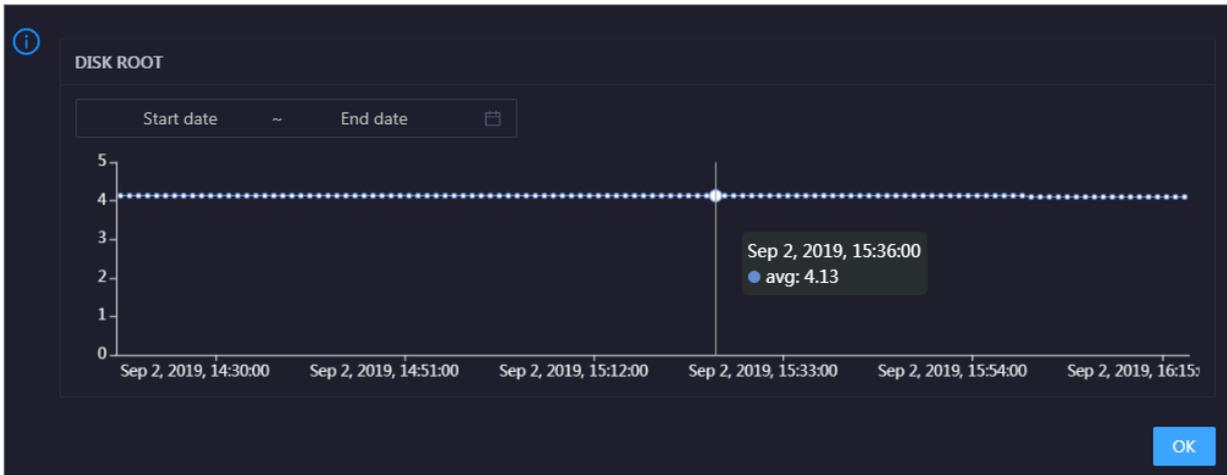


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

This chart displays the trend line of the average root disk usage (avg) for the selected service over time.

Click  in the upper-right corner of the chart to zoom in it.

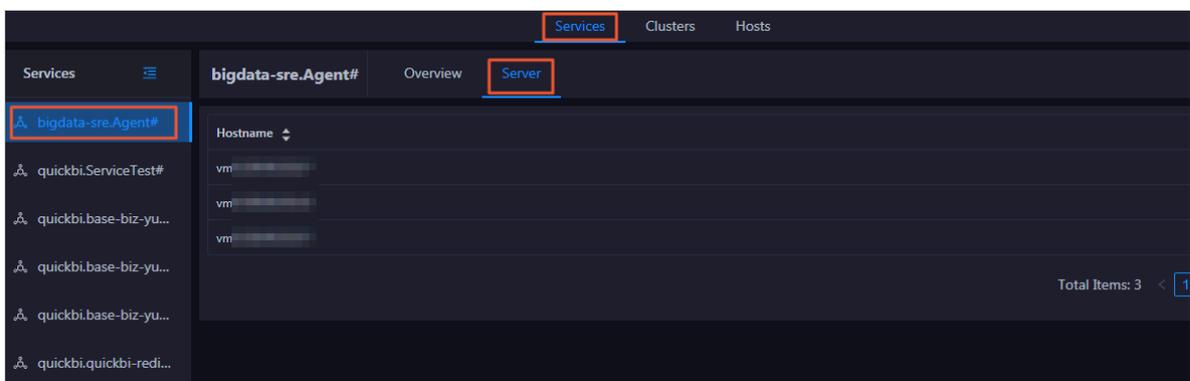


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

3.1.11.2.2 Service hosts

Apsara Bigdata Manager (ABM) allows you to view the host list of each Quick BI service so that you can understand the service deployment on hosts.

1. At the top of the O&M page, click **Services**.
2. On the **Services** page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.
3. Click the **Server** tab. The **Server** page for the service appears.



On the **Server** page, you can view the hosts where the selected service is run.

3.1.11.3 Cluster O&M

3.1.11.3.1 Cluster overview

The cluster overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.

Entry

1. At the top of the O&M page, click Clusters.
2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.



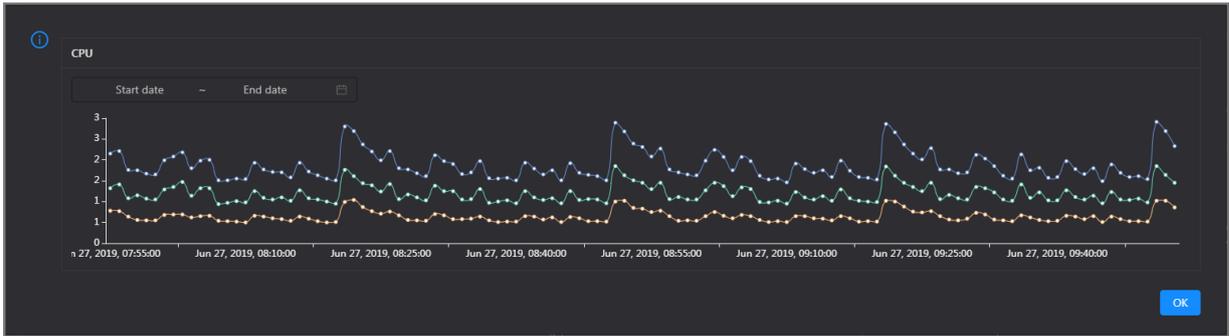
The Overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the cluster. To view information about a cluster, select a region in the left-side navigation pane, and then select a cluster in the region.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

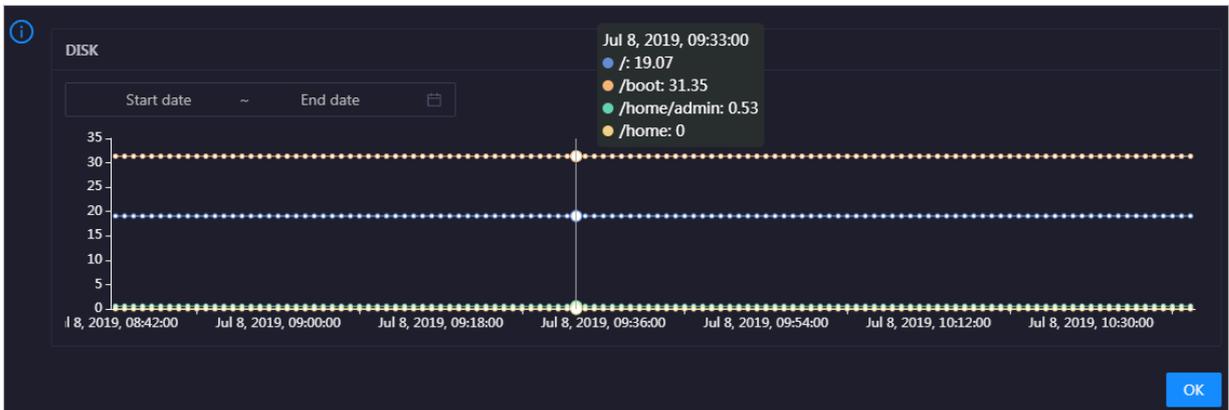
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the /, /boot, /home/admin, and /home directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

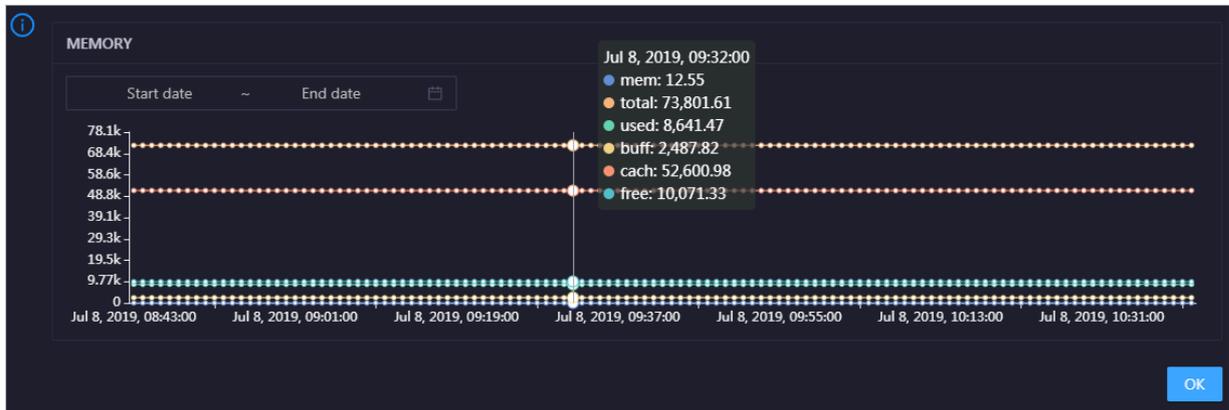


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

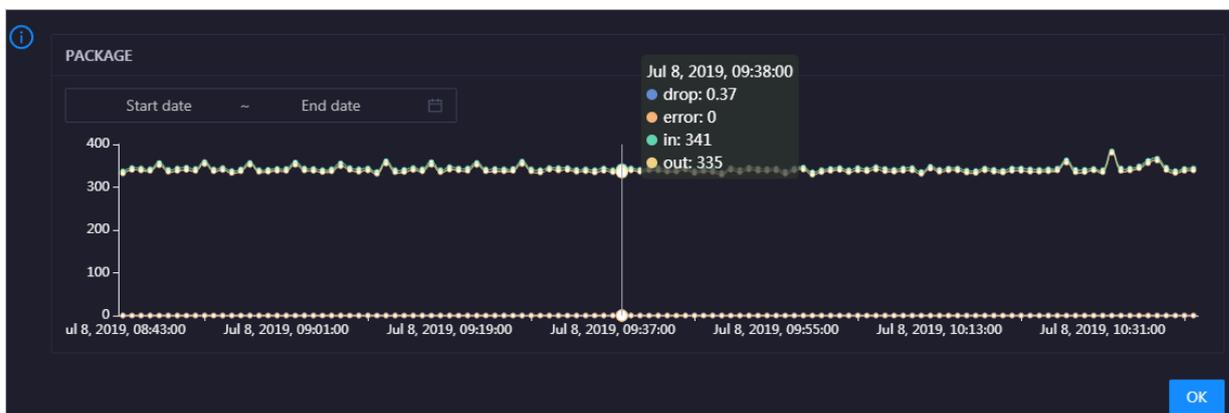


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

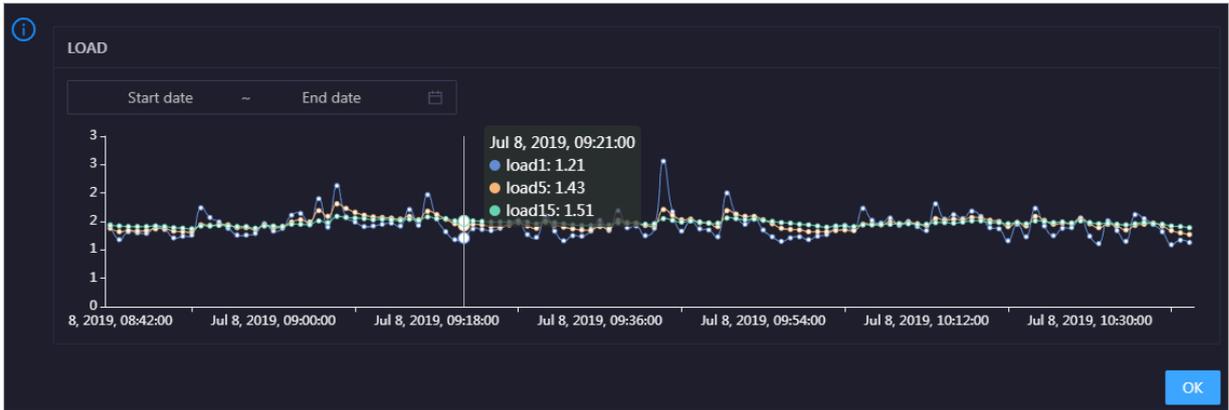


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

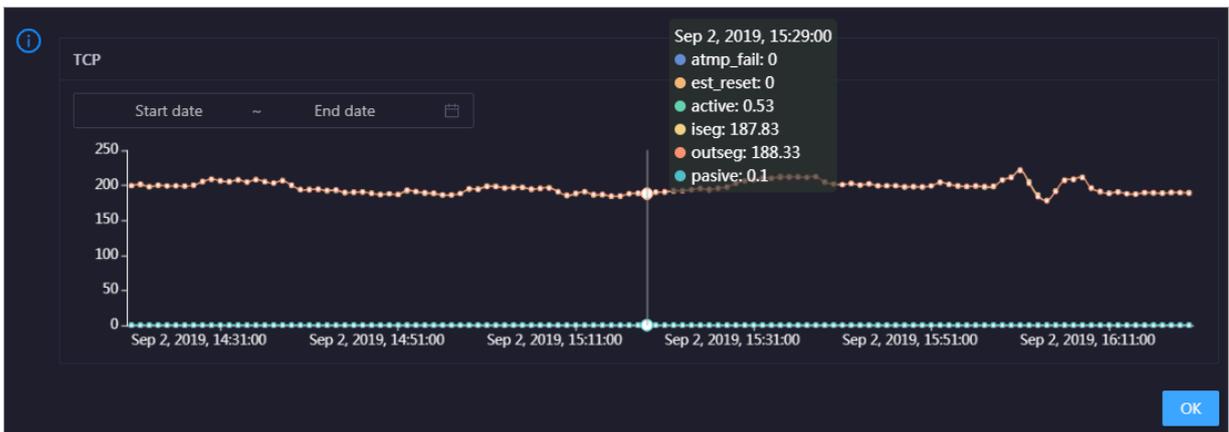


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

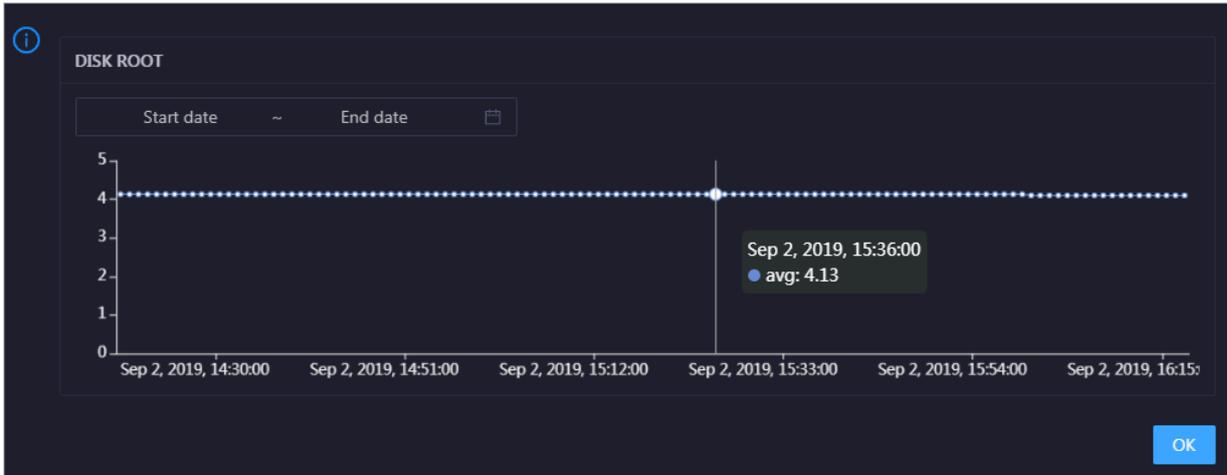


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

3.1.11.3.2 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. At the top of the O&M page, click Clusters.
2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.

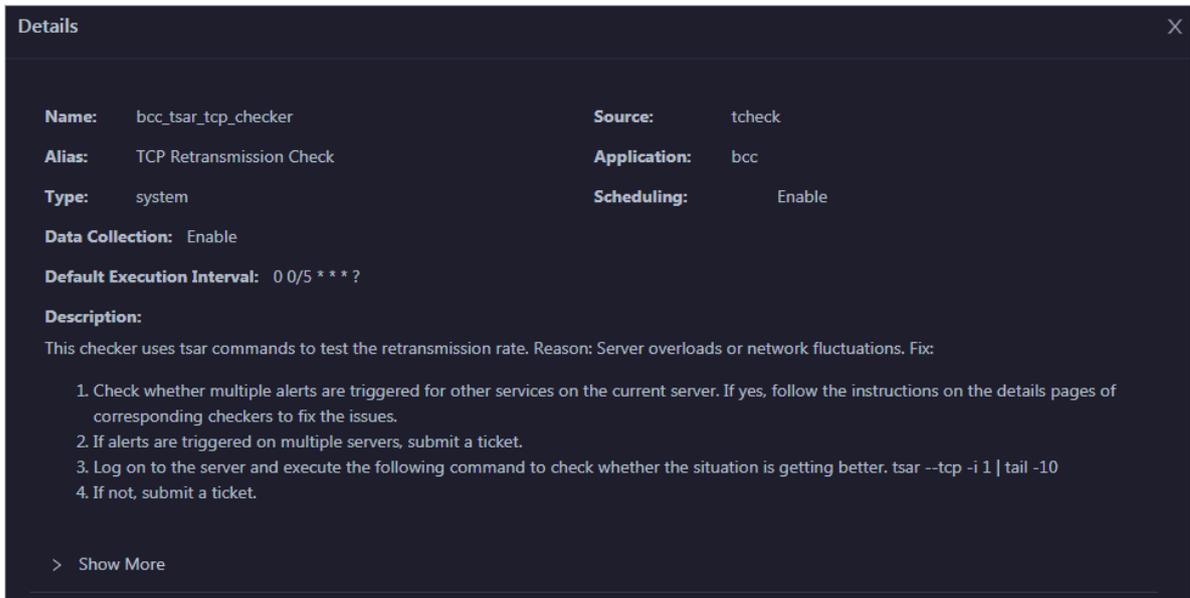
Checker	Source	Critical	Warning	Exception	Actions
+ bcc_check_ntp	tcheck	0	3	0	Details
+ bcc_tsar_top_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_network_top_connections_checker	tcheck	0	0	0	Details
+ bcc_disk_usage_checker	tcheck	0	0	0	Details
+ bcc_host_live_check	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among

them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

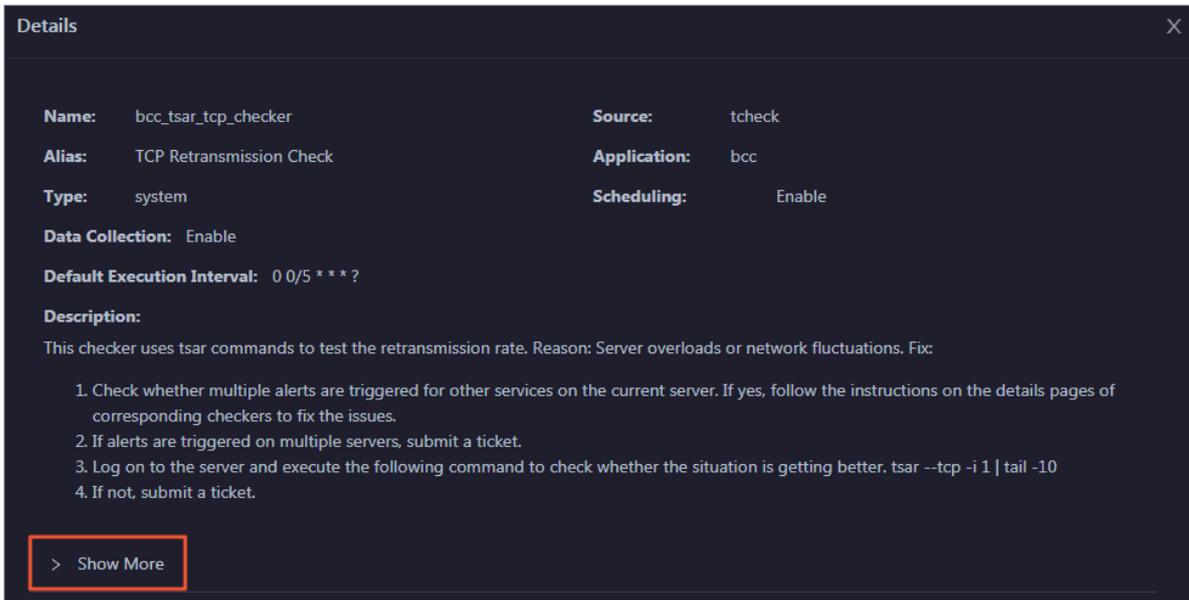
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

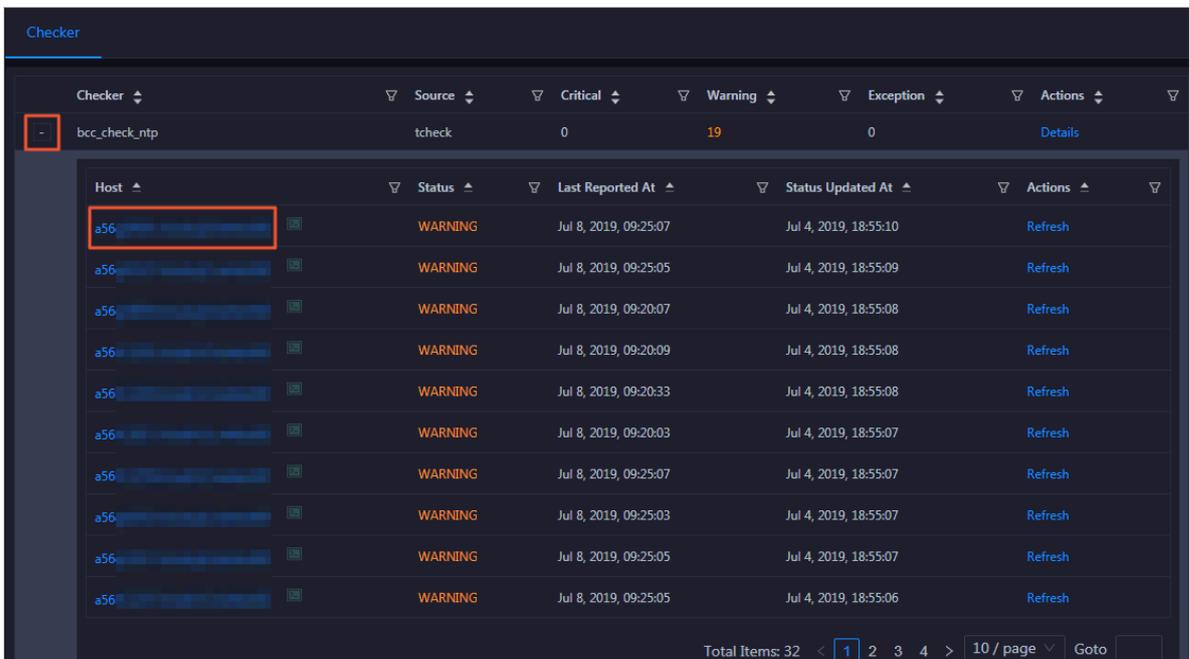


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

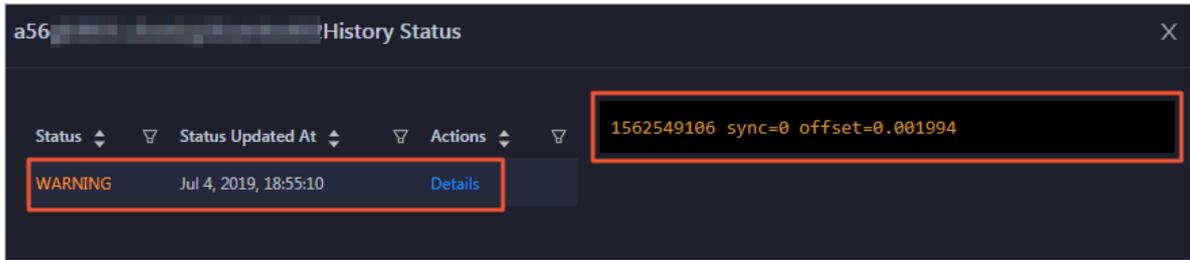
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.

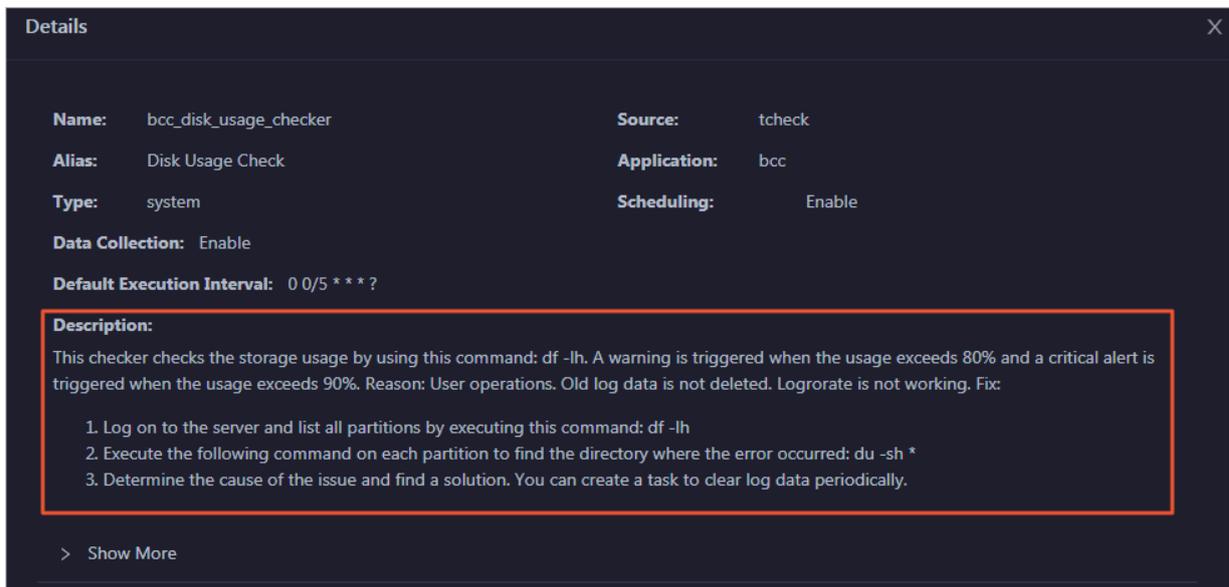


2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

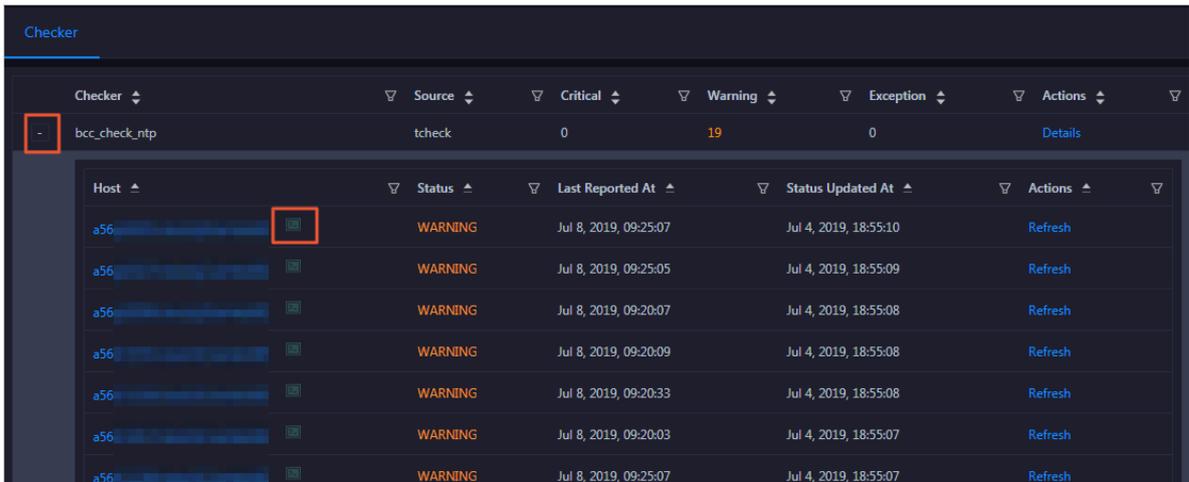
- On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



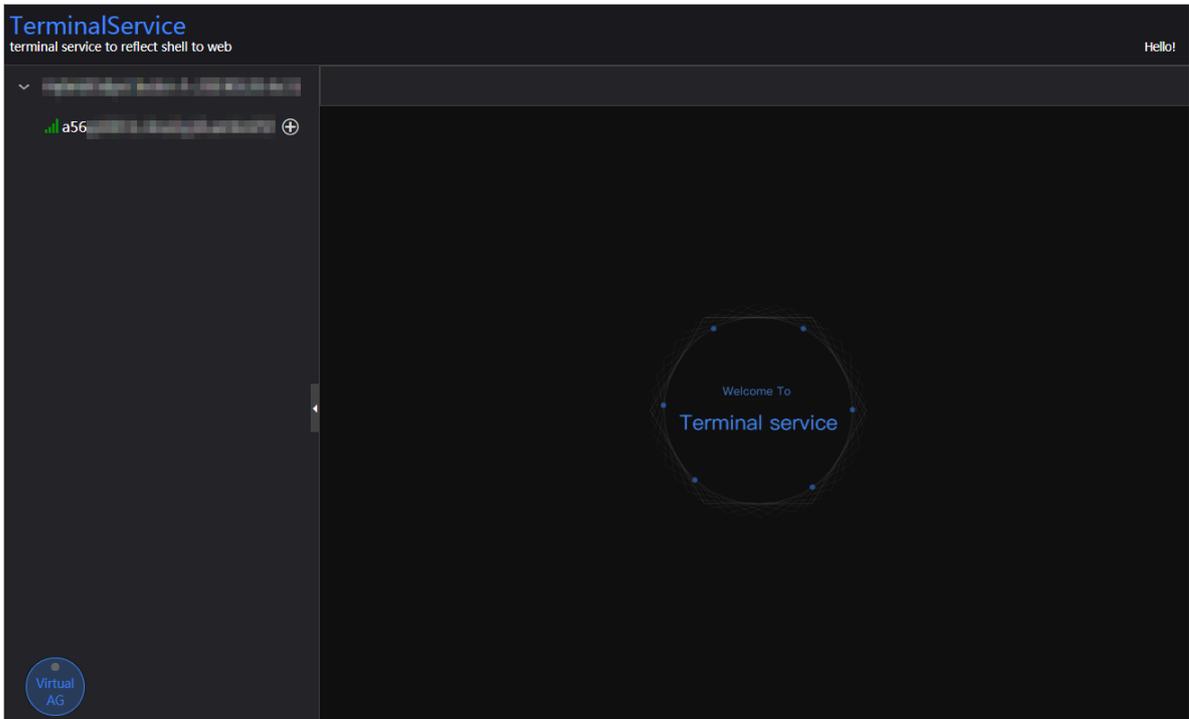
Log on to a host

- To log on to a host to clear alerts or perform other operations, follow these steps:

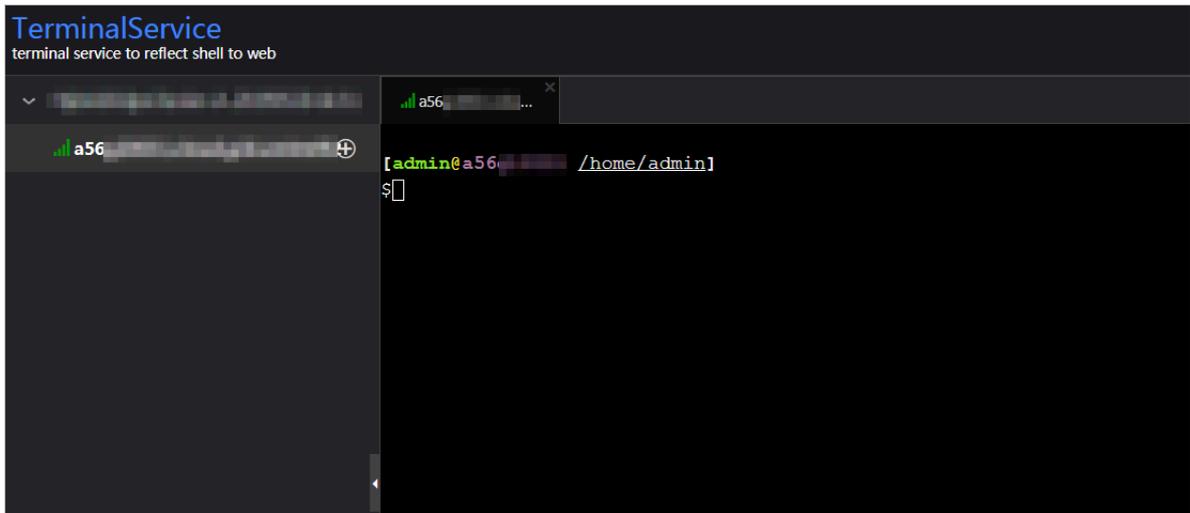
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.



3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click **Refresh** in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

Checker	Source	Critical	Warning	Exception	Actions
bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

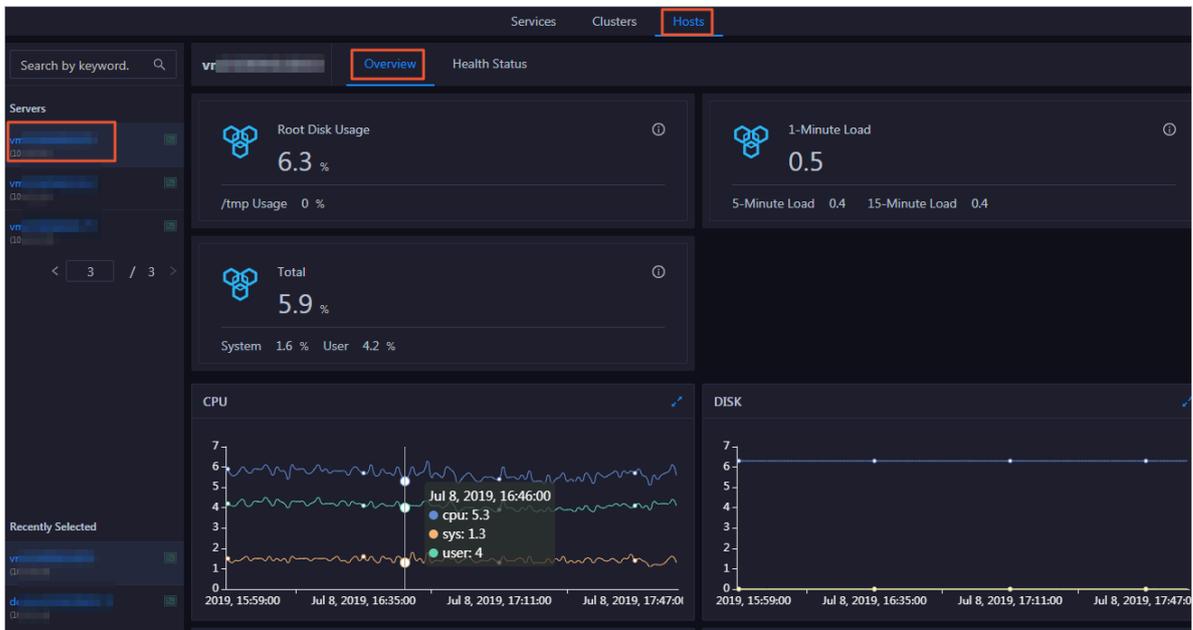
3.1.11.4 Host O&M

3.1.11.4.1 Host overview

The host overview page displays the overall running information about a host in an ElasticSearch cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Entry

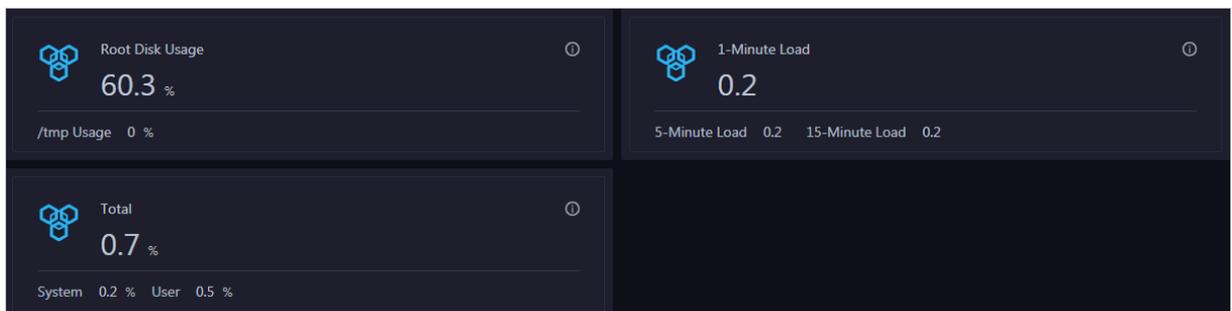
1. At the top of the O&M page, click Hosts.
2. On the Hosts page, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.



On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Root Disk Usage, Total, and 1-Minute Load

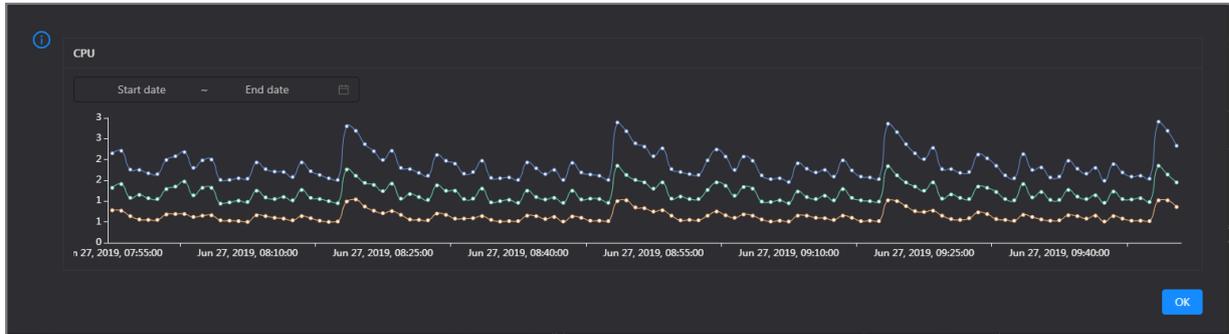
These sections display the root disk usage, total usage, and 1-minute load for the selected host. The Root Disk Usage section provides the usage of the `/tmp` directory. The Total section provides the system usage and user usage. The 1-Minute Load section provides the 1-minute, 5-minute, and 15-minute load averages.



CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

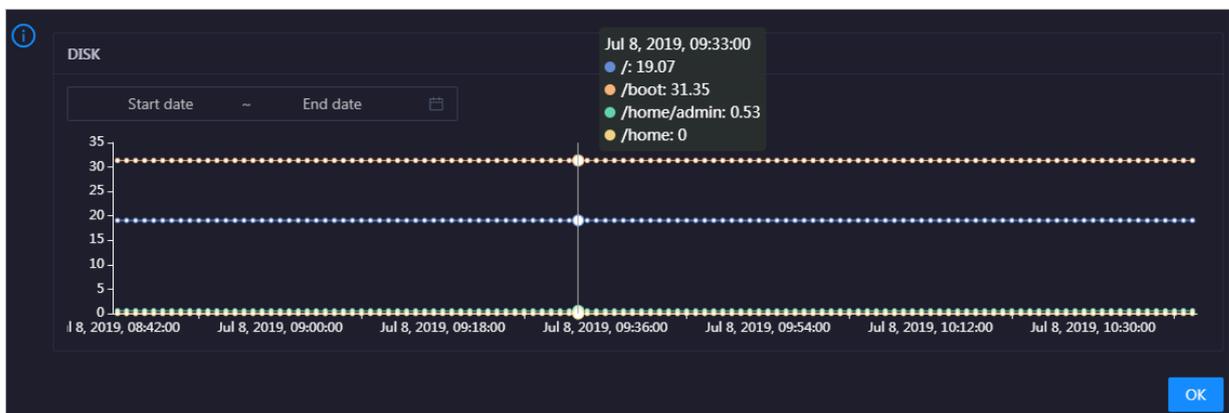


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



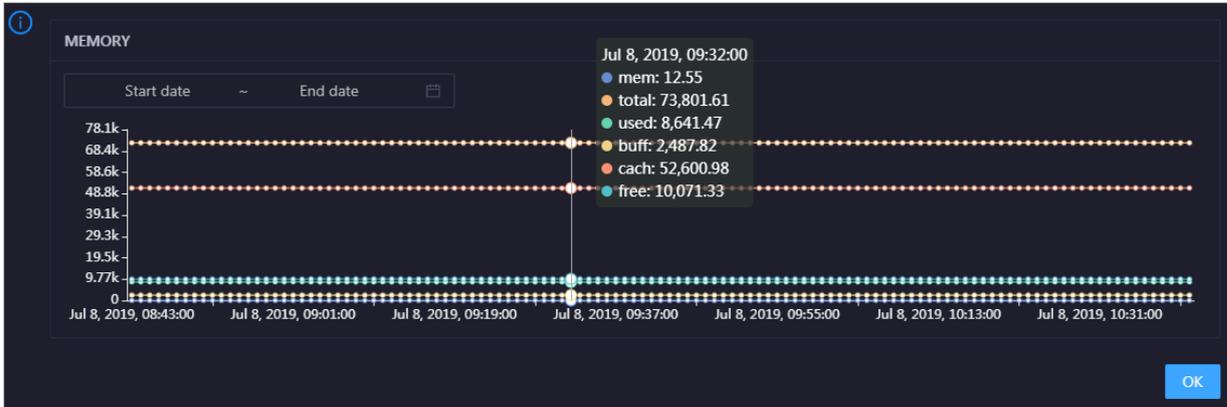
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size

of memory used by the page cache (cach), and available memory size (free) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

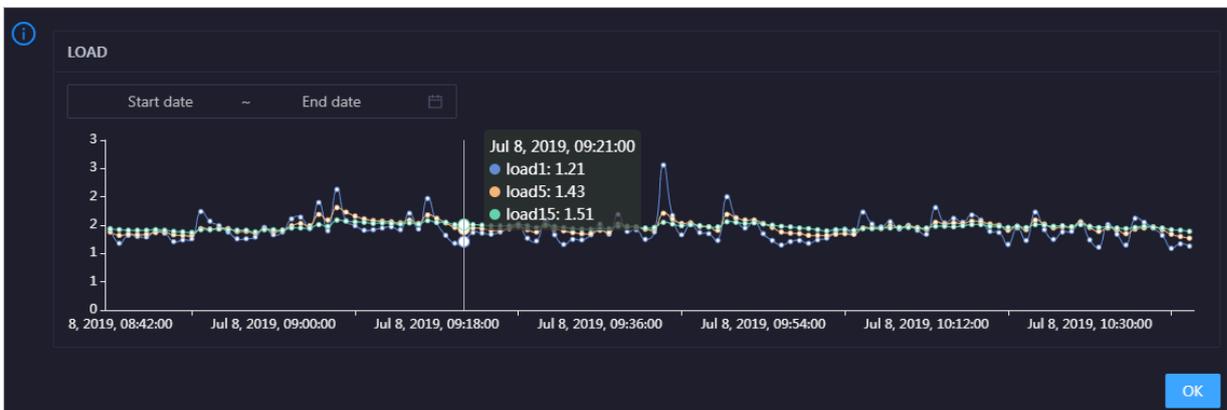


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



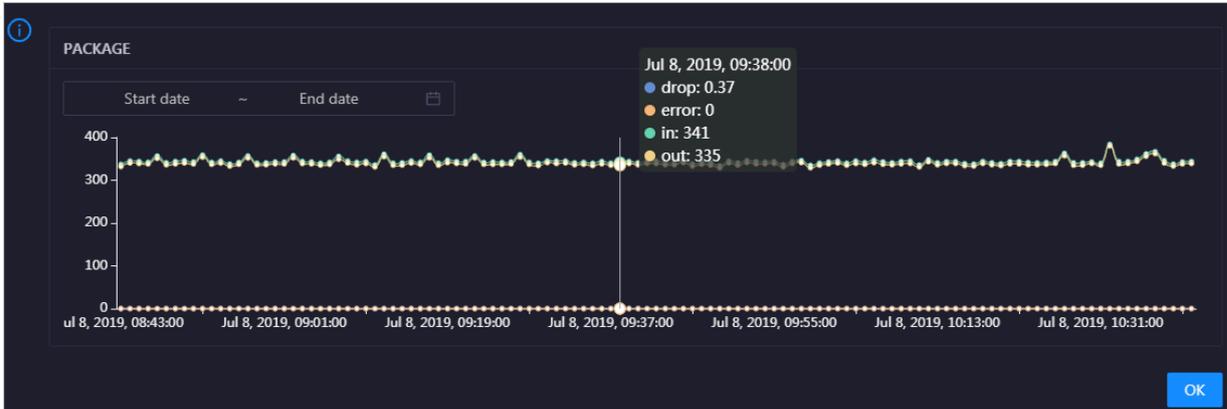
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for

the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.

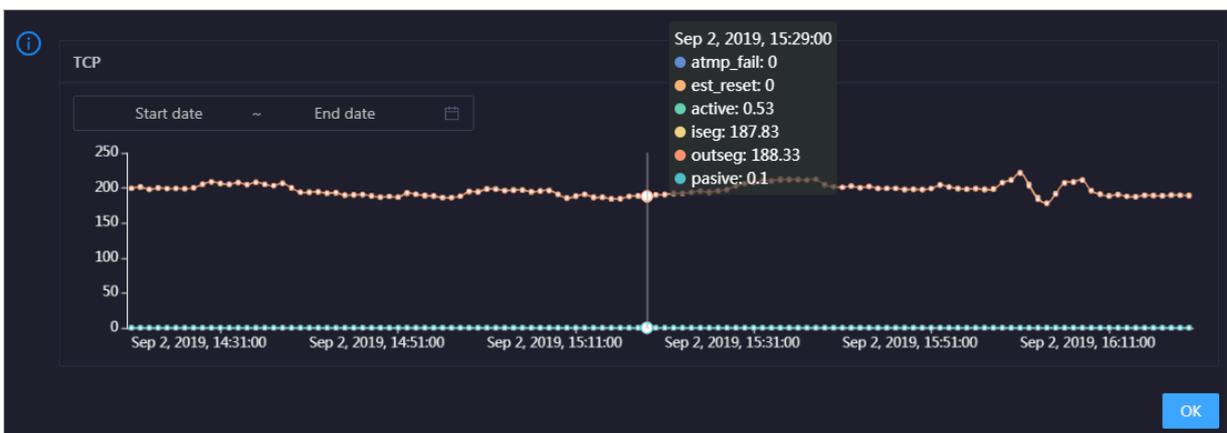


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click  in the upper-right corner of the chart to zoom in the chart.

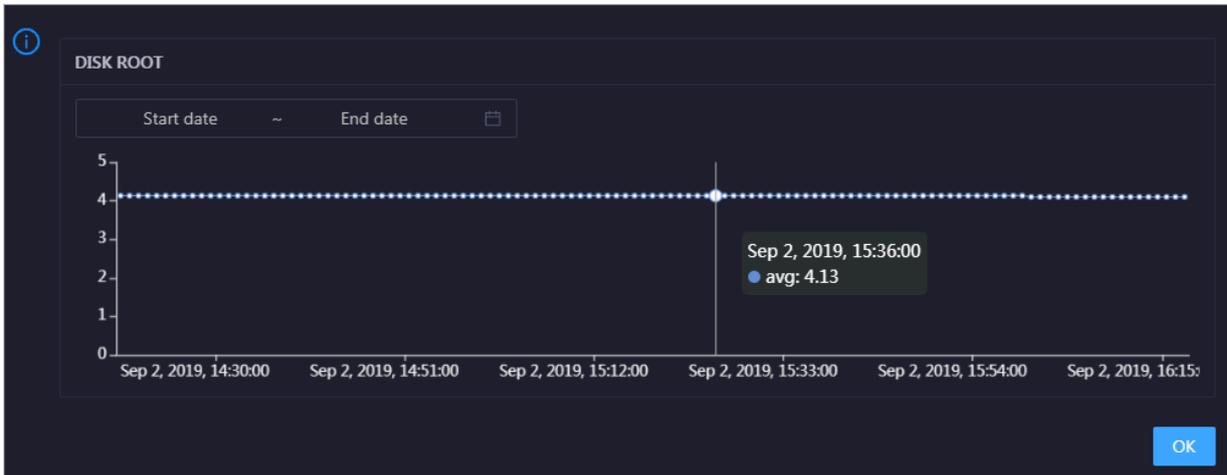


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the host over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

Health Check [View Details](#)

Currently, 9 checkers are deployed on the service. 2 critical, 0 exception, and 0 warning alerts are reported.

Click [View Details](#) to go to the [Host health](#) page. On this page, you can view the health check details.

Health Check History

This section displays a record of the health checks performed on the host.

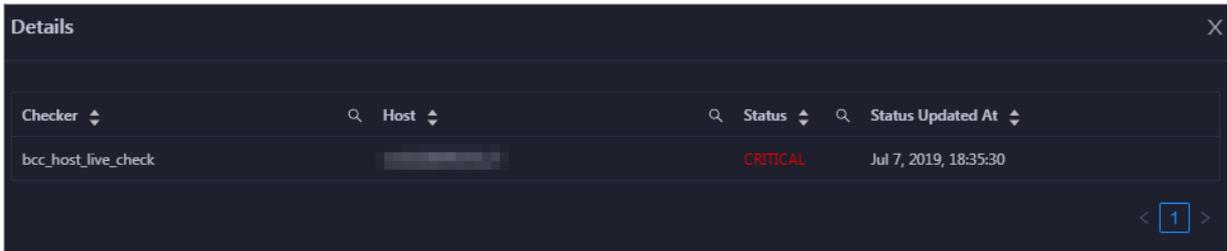
Health Check History [View Details](#)

Time	Event Content
Recently	1 alerts are reported by checkers.

< 1 >

Click [View Details](#) to go to the [Host health](#) page. On this page, you can view the health check details.

You can click the event content of a check to view the exception items.

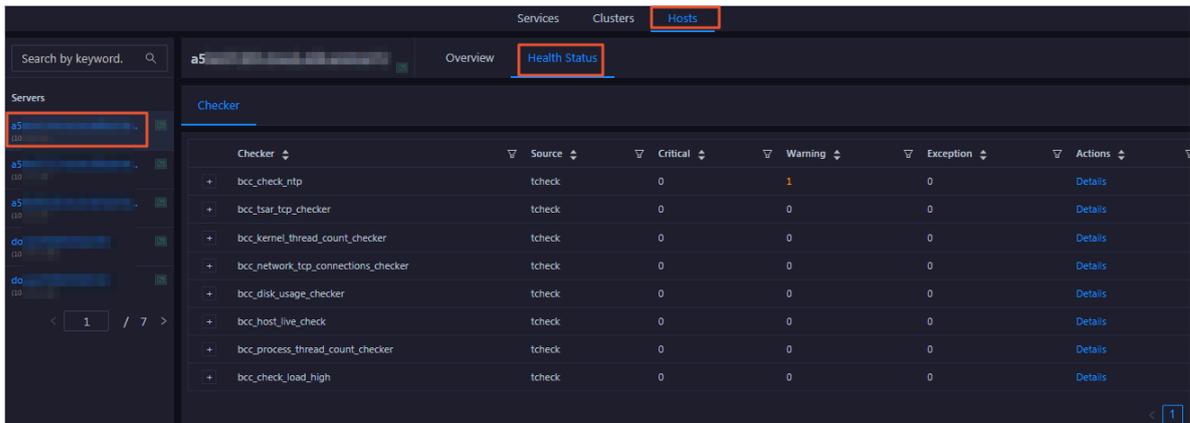


3.1.11.4.2 Host health

On the host health status page, you can view the checkers of all hosts, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

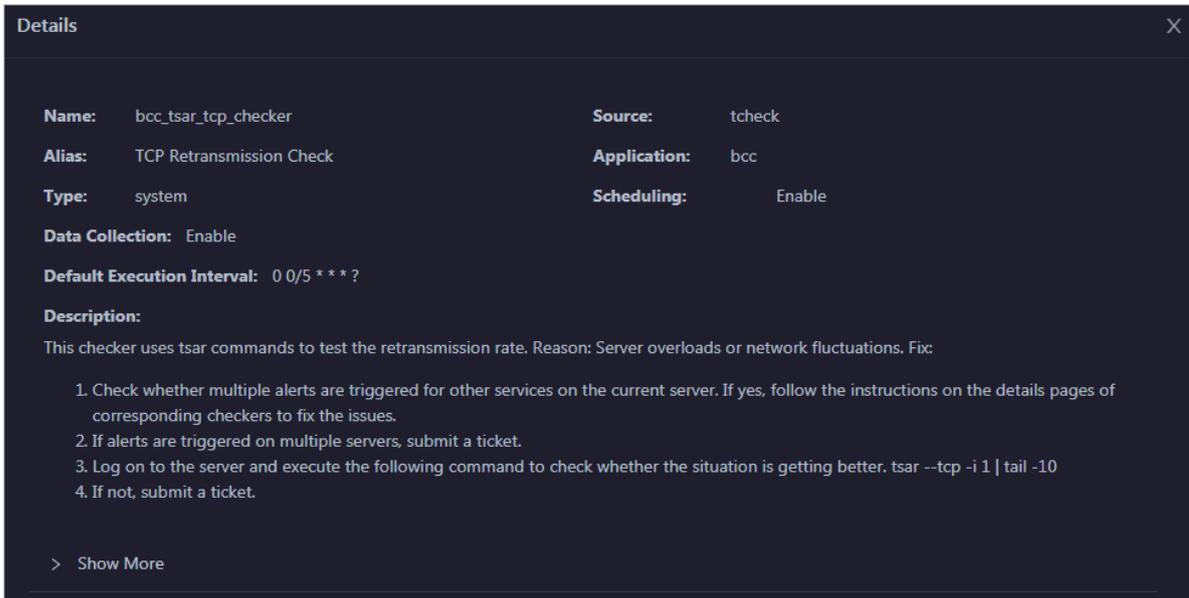
1. At the top of the O&M page, click Hosts.
2. On the Hosts page, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.



On the Health Status page, you can view all checkers of the host and the check results for the hosts in the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

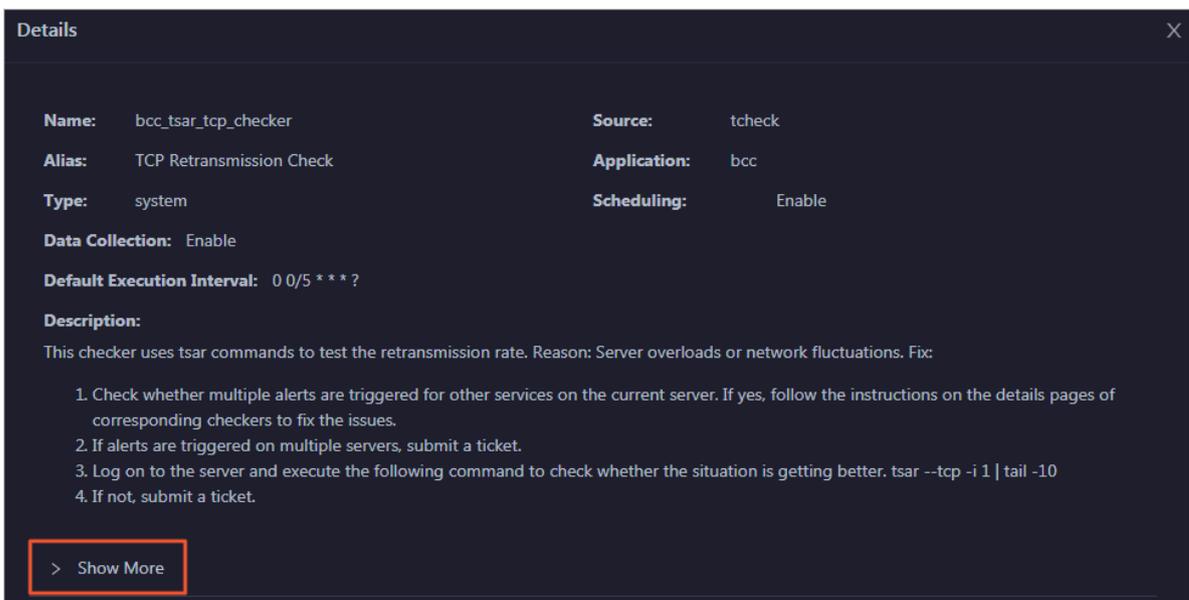
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

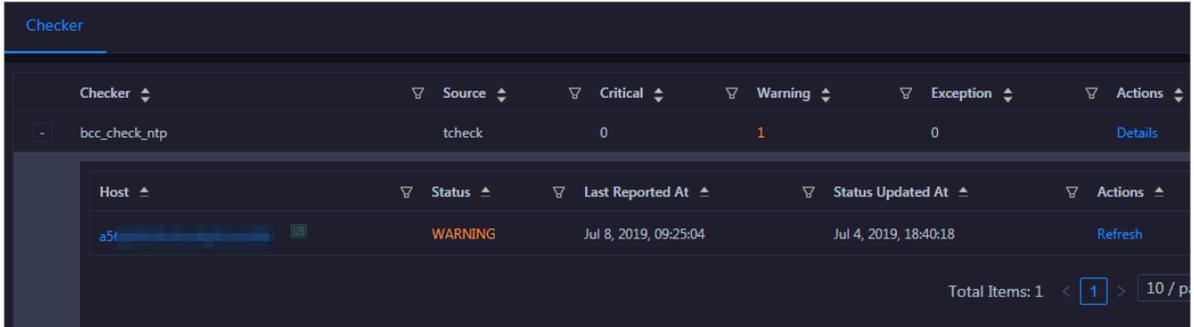


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

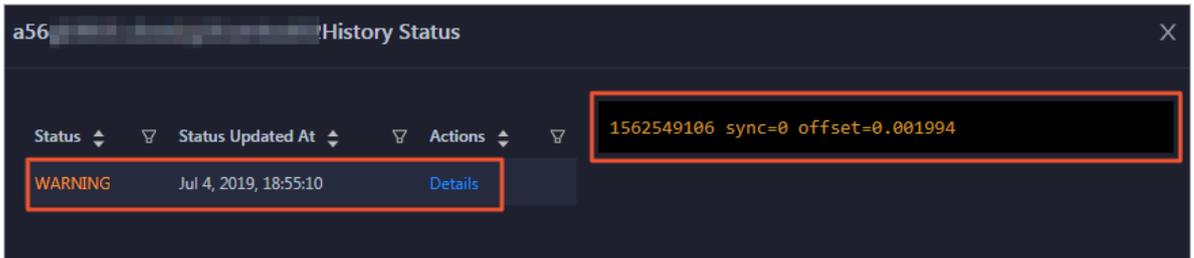
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

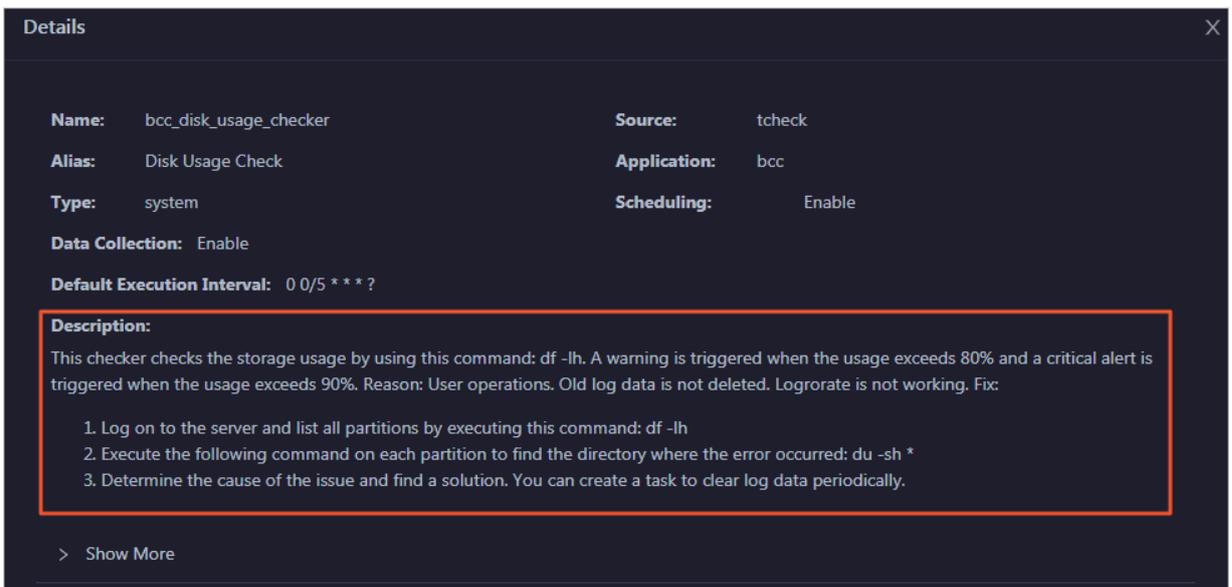


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

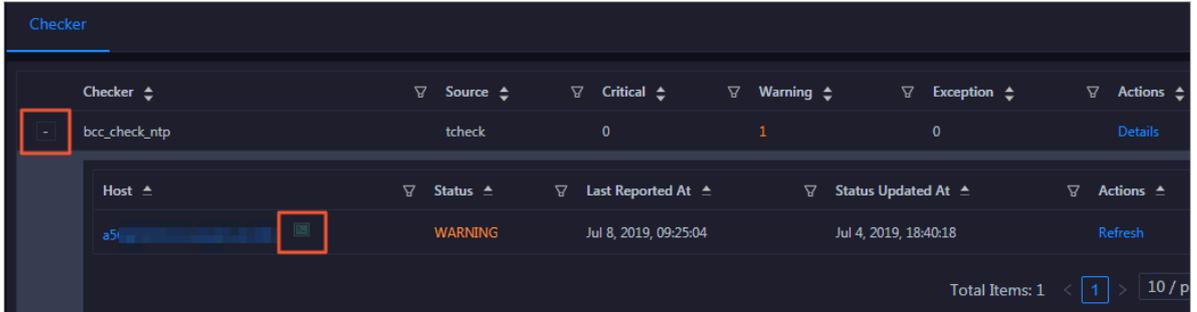
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



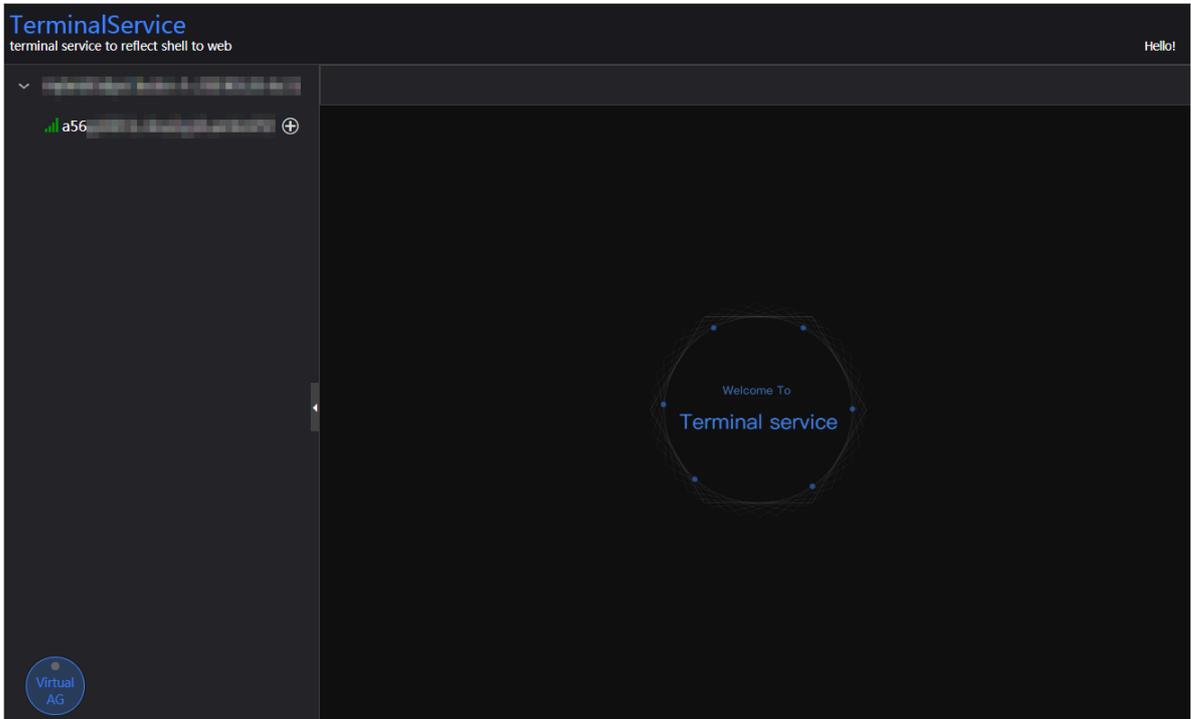
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

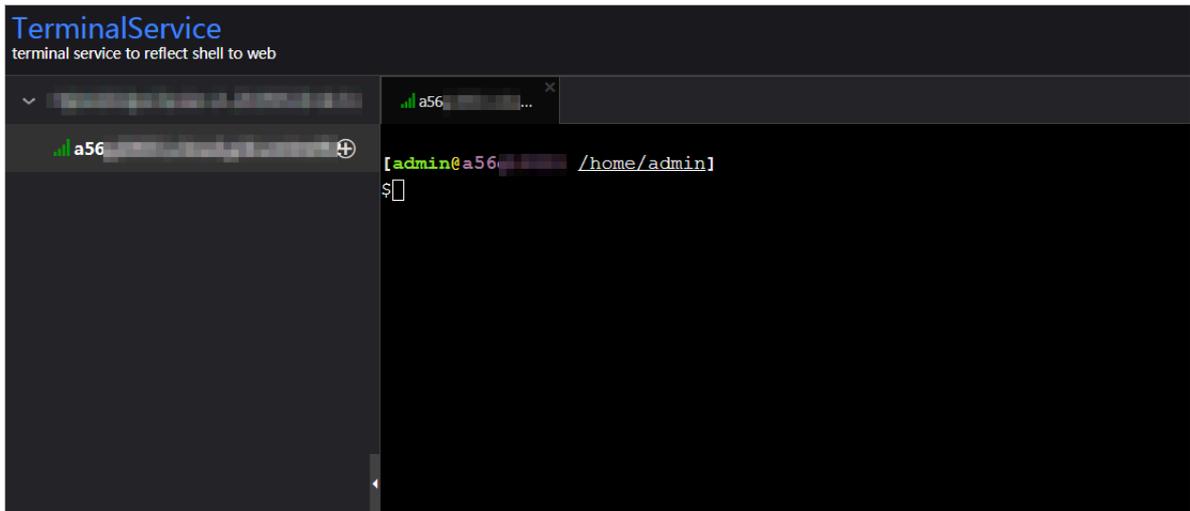
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

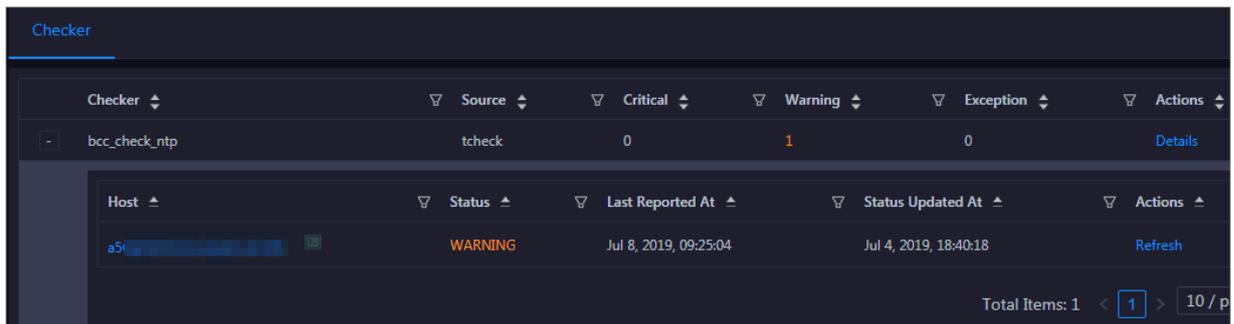


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click **Refresh** in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.12 PAI

3.1.12.1 O&M overview

This topic describes the features of PAI O&M and how to access the PAI O&M page.

Modules

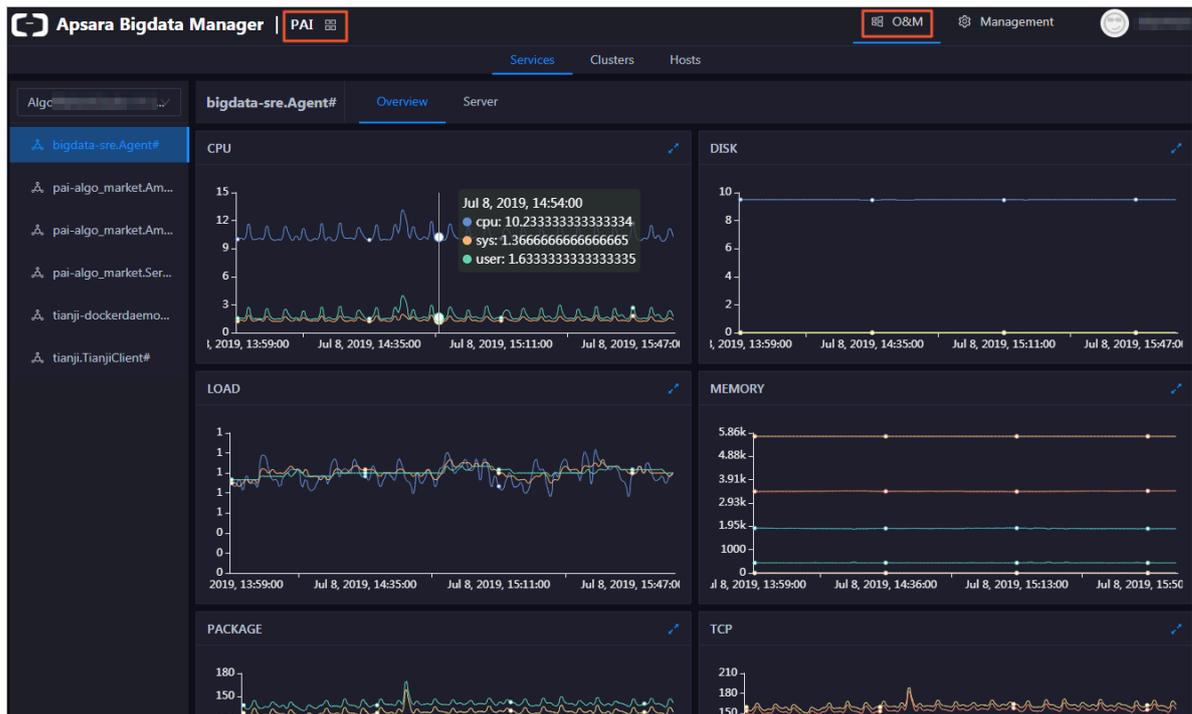
PAI O&M includes service O&M, cluster O&M, and host O&M. The following table describes them in detail.

Module	Feature	Description
Service O&M	Service overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service in a cluster.
	Service hosts	Displays the host list of each service in a cluster so that you can understand the service deployment on hosts.
Cluster O&M	Cluster overview	Displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.
	Cluster health	Displays the check results for a cluster. The check results are divided into the Critical, Warning, Exception, and OK types.
Host O&M	Host overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Host health	Displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.

Entry

1. [Log on to the ABM console.](#)
2. Click  in the upper-left corner, and then click PAI.

3. On the page that appears, click O&M at the top. The Services page appears.



The O&M page includes three modules, namely, Services, Clusters, and Hosts.

3.1.12.2 Service O&M

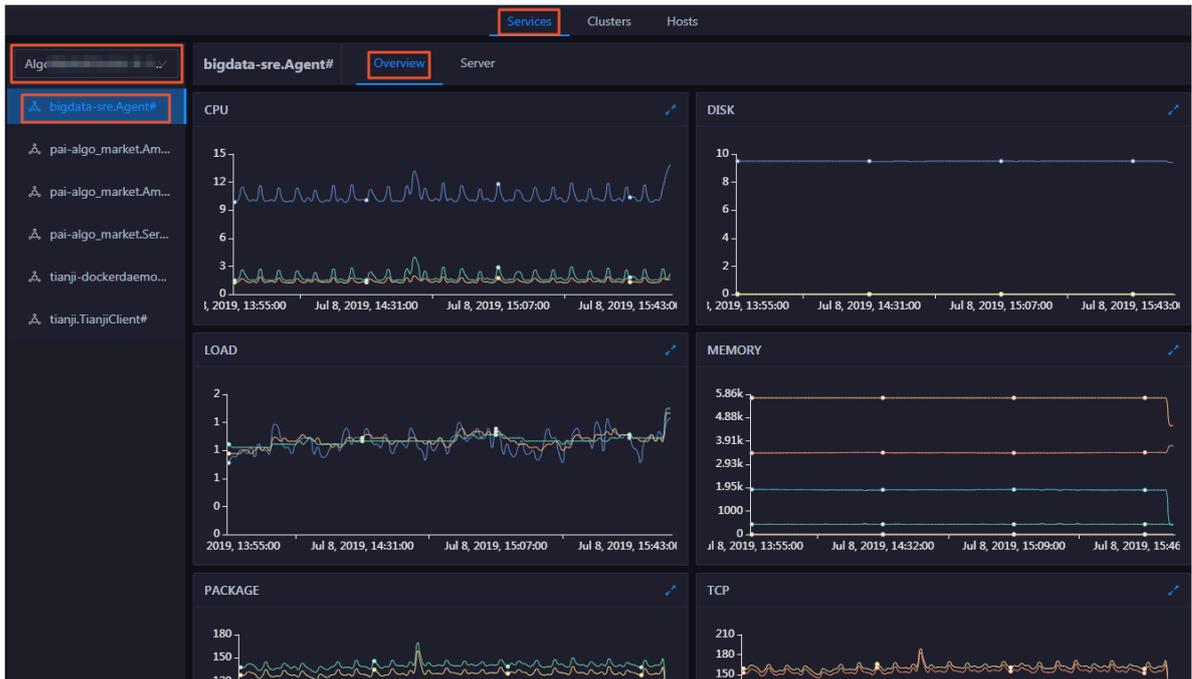
3.1.12.2.1 Service overview

The service overview page lists all PAI services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.

Entry

1. At the top of the O&M page, click Services.
2. On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.

3. Click the Overview tab. The Overview page for the service appears.



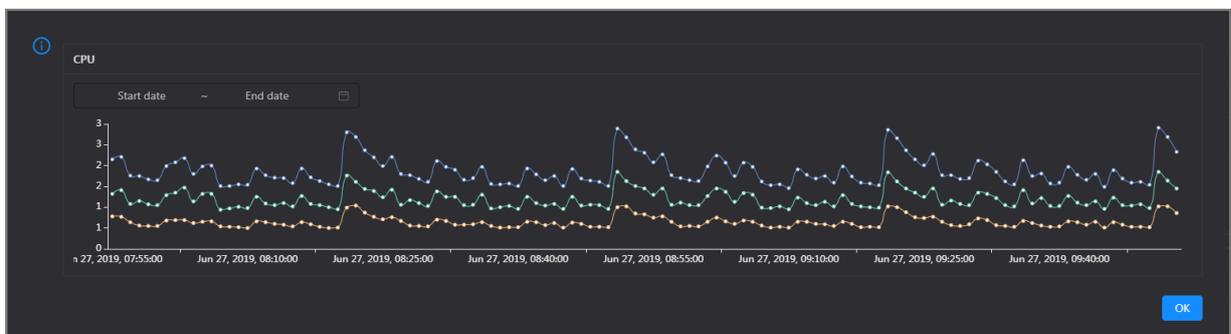
On the Overview page, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

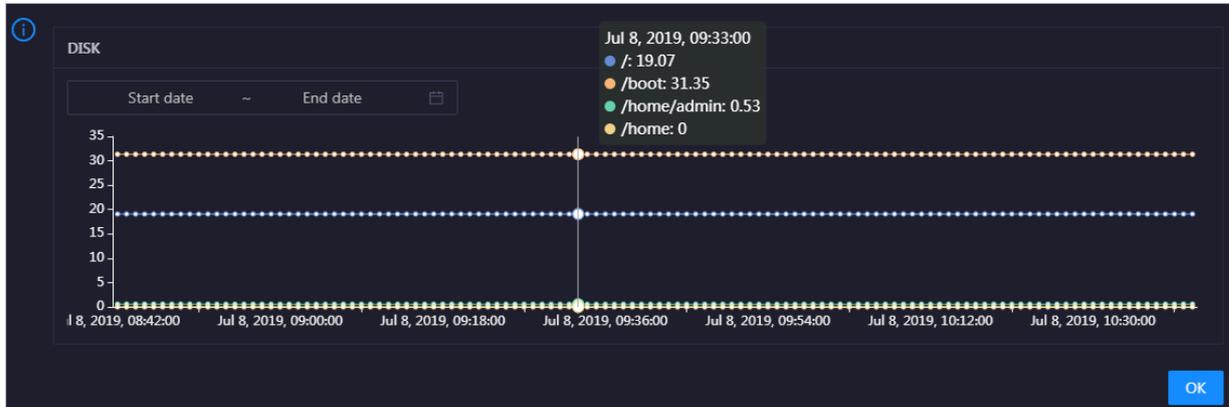
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the service in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

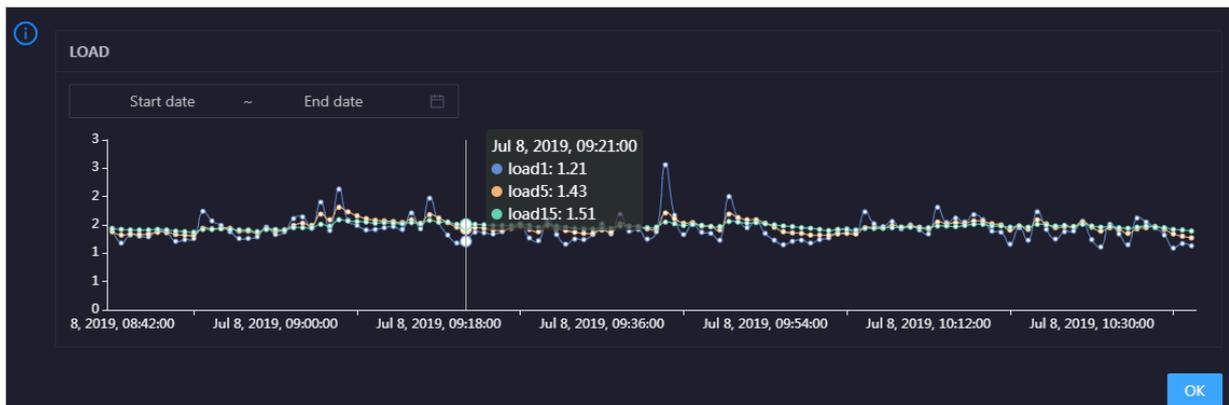


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

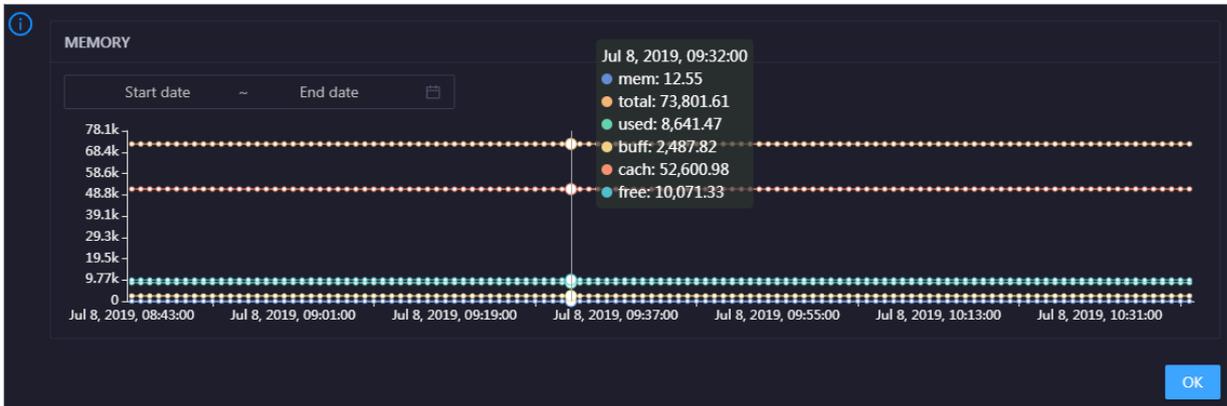


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

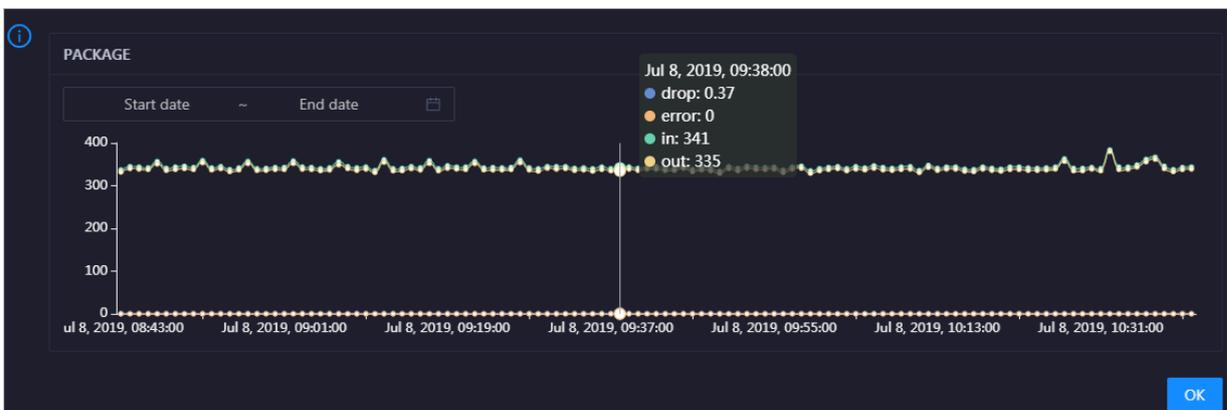


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the selected service over time in different colors. These trend lines reflect the data transmission status of the service.

Click  in the upper-right corner of the chart to zoom in it.

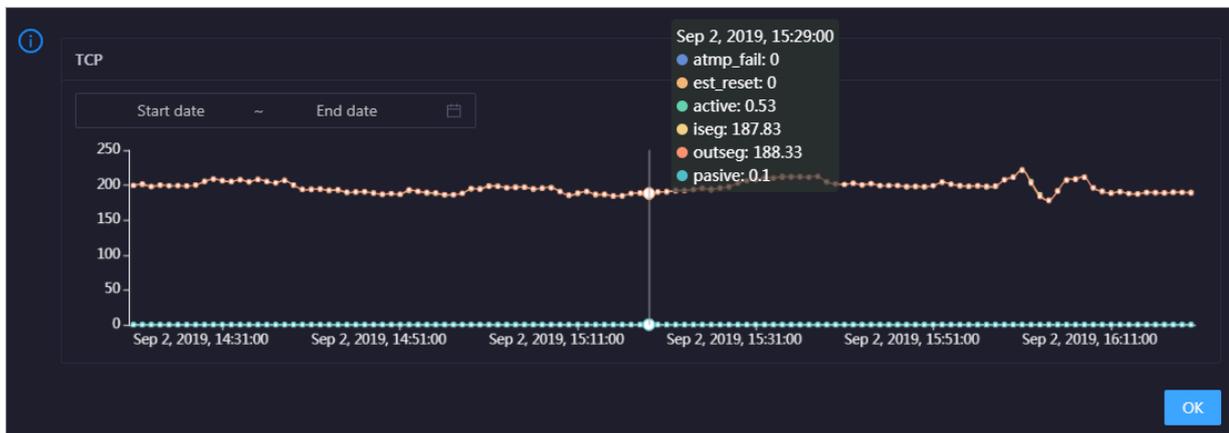


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the selected service over time in different colors. These trend lines reflect the TCP connection status of the service.

Click  in the upper-right corner of the chart to zoom in it.

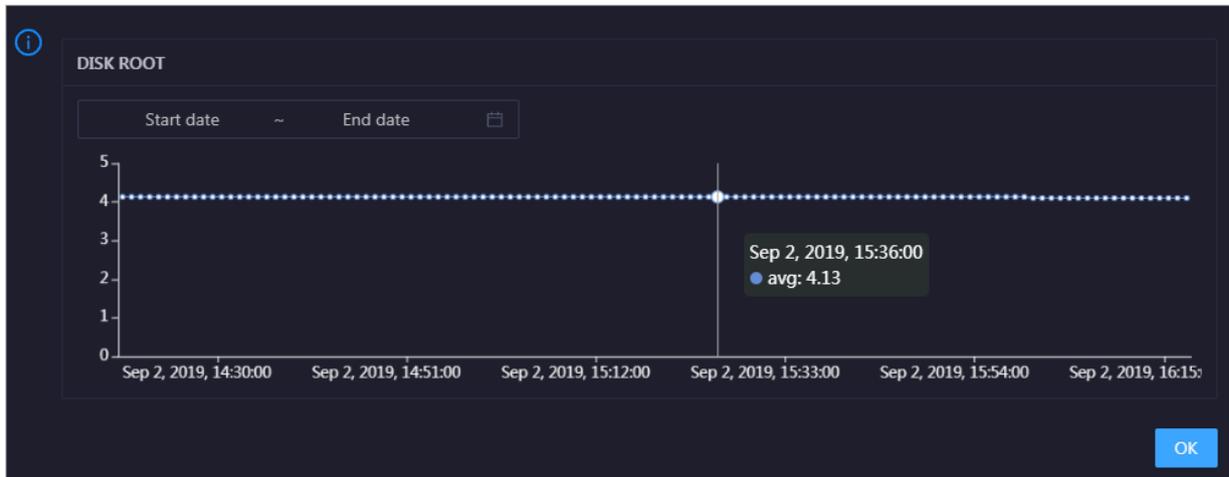


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

This chart displays the trend line of the average root disk usage (avg) for the selected service over time.

Click  in the upper-right corner of the chart to zoom in it.

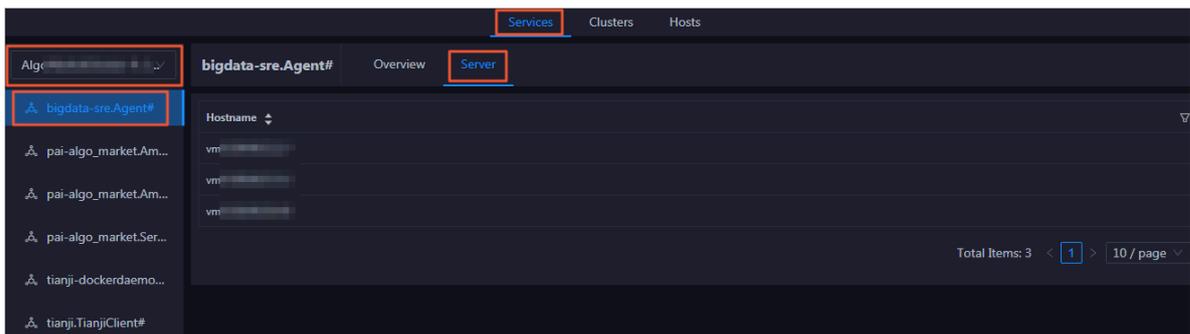


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

3.1.12.2.2 Service hosts

Apsara Bigdata Manager (ABM) allows you to view the host list of each PAI service so that you can understand the service deployment on hosts.

1. At the top of the O&M page, click Services.
2. On the Services page, search for a cluster in the search box above the left-side service list, and then select a service in the service list.
3. Click the Server tab. The Server page for the service appears.



On the Server page, you can view the hosts where the selected service is run.

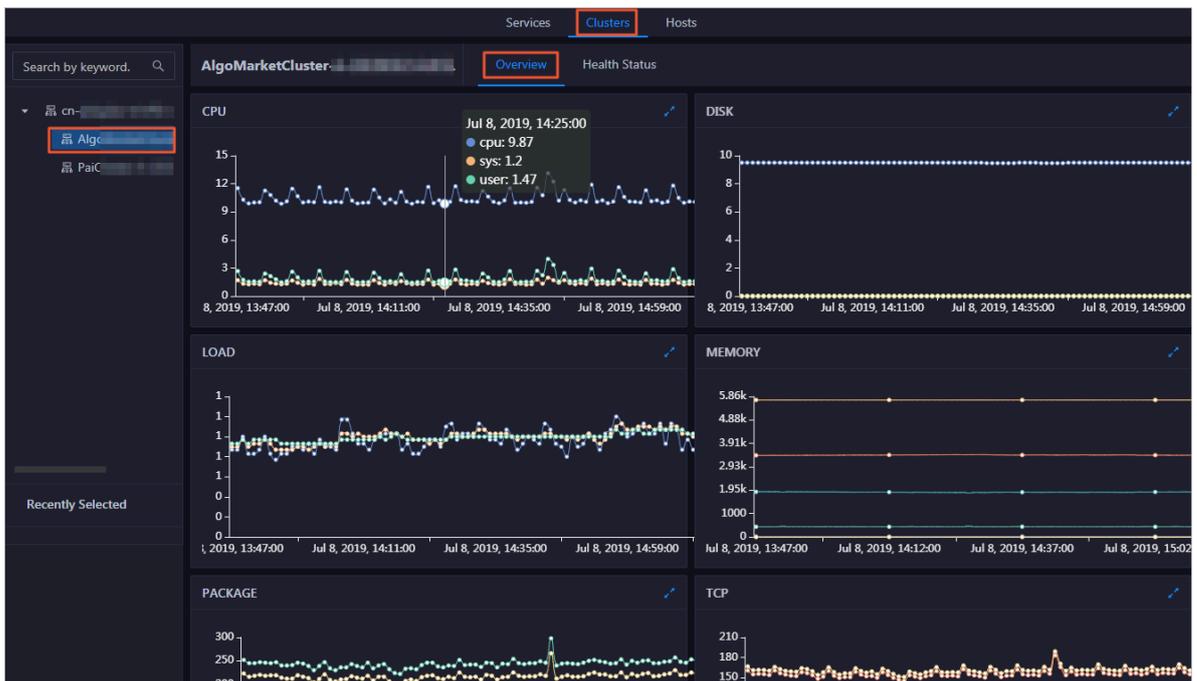
3.1.12.3 Cluster O&M

3.1.12.3.1 Cluster overview

The cluster overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for a cluster.

Entry

1. At the top of the O&M page, click Clusters.
2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.



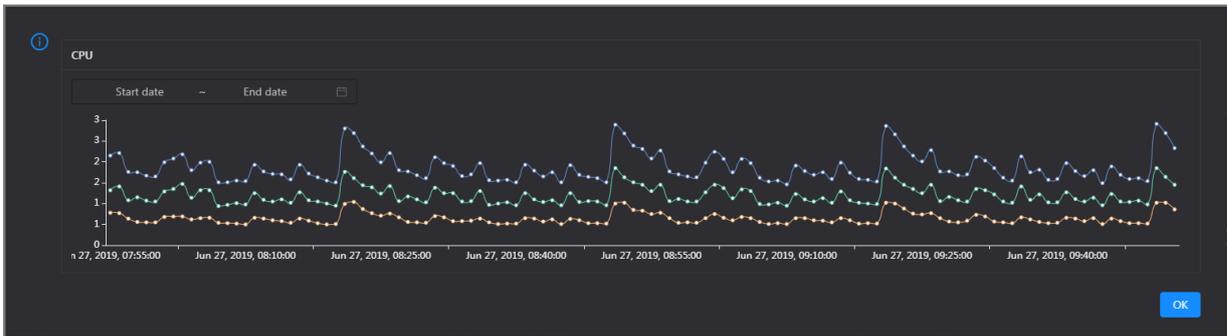
The Overview page displays the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the cluster. To view information about a cluster, select a region in the left-side navigation pane, and then select a cluster in the region.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

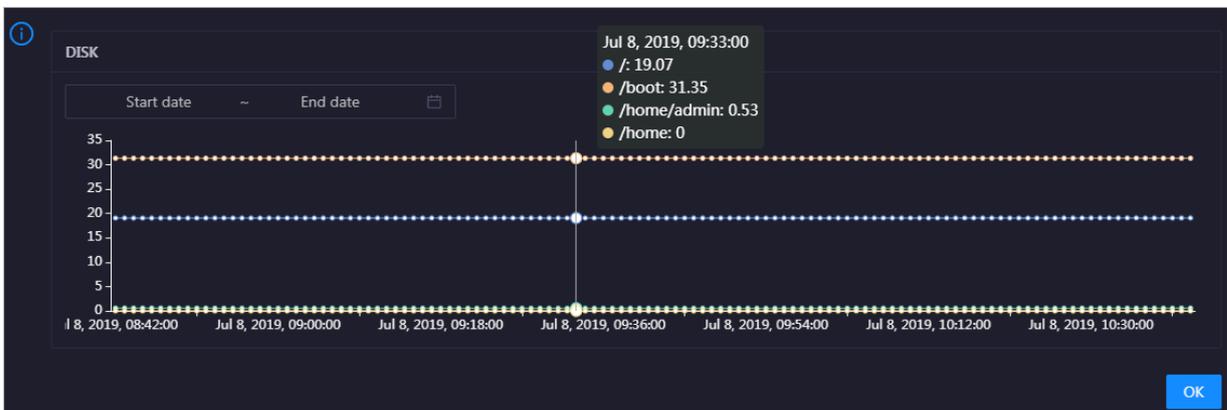
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

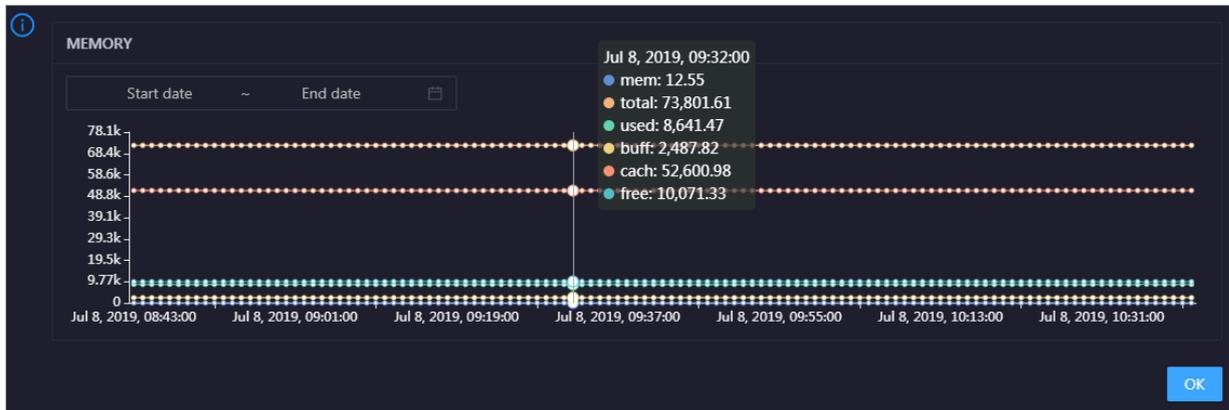


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (`mem`), total memory size (`total`), used memory size (`used`), size of memory used by kernel buffers (`buff`), size of memory used by the page cache (`cach`), and available memory size (`free`) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

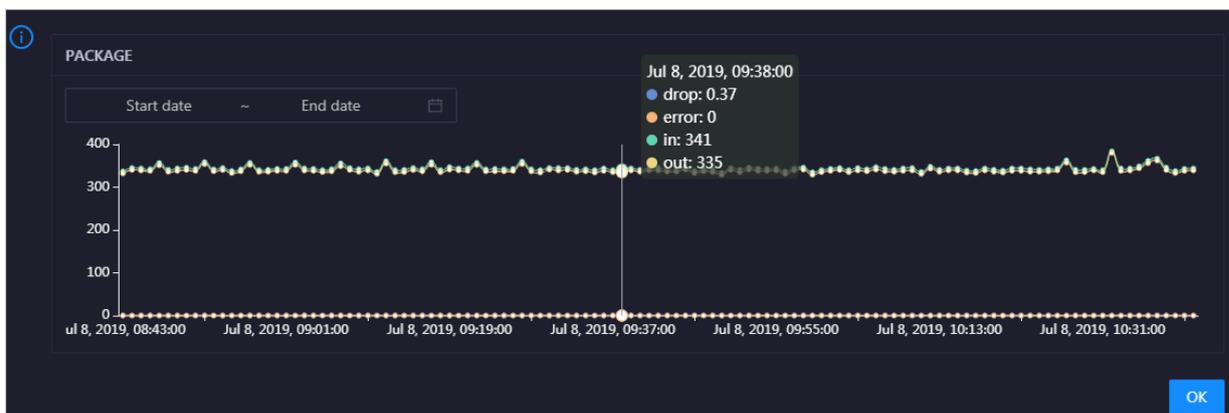


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

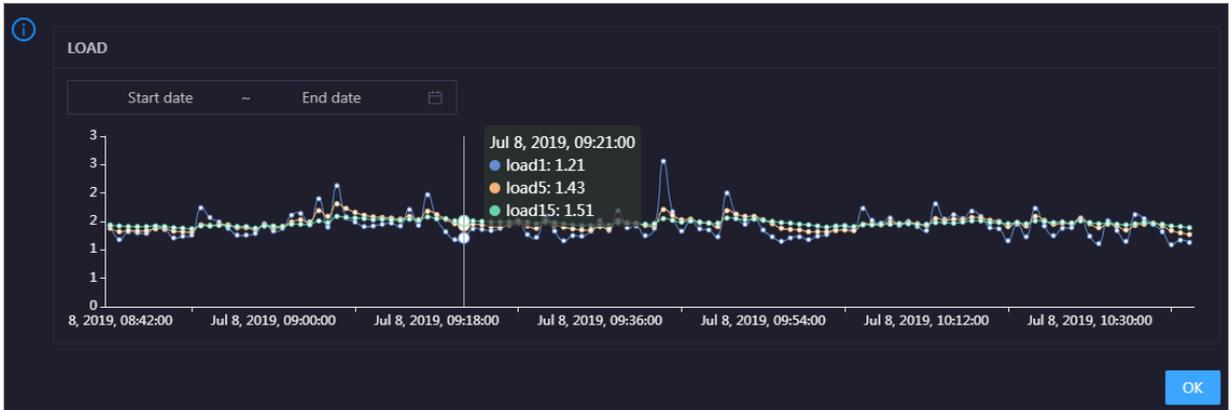


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

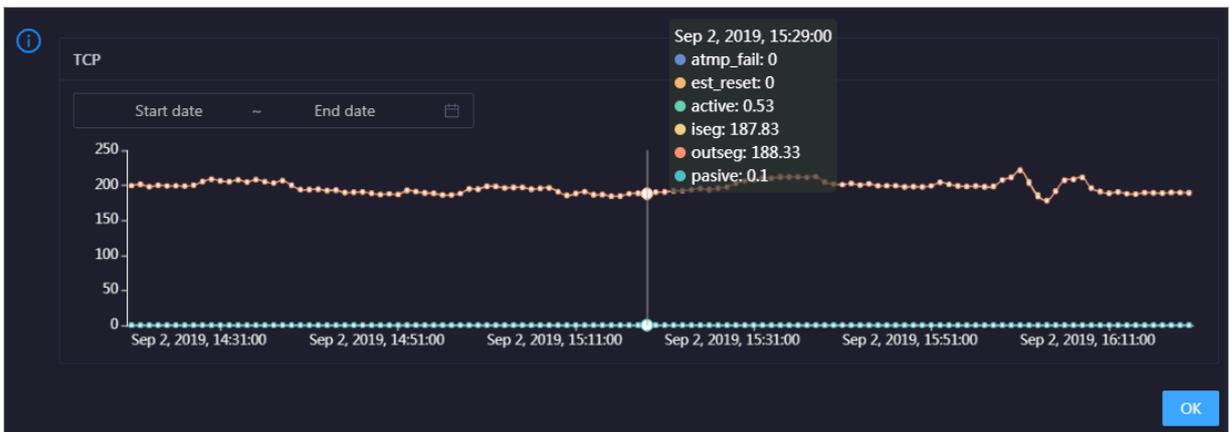


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

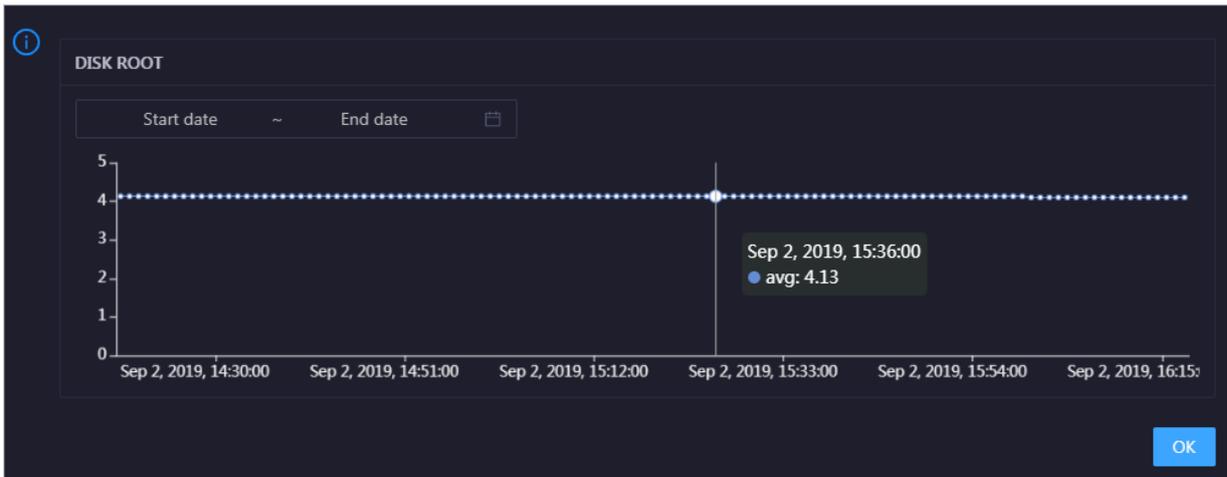


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

3.1.12.3.2 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

1. At the top of the O&M page, click Clusters.
2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.

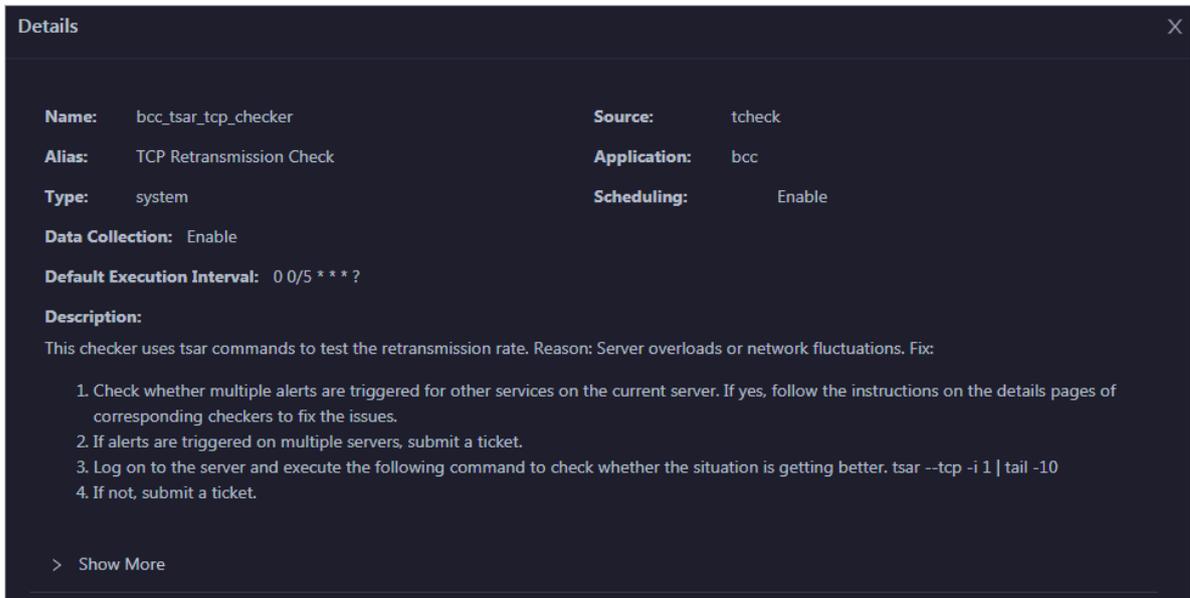
Checker	Source	Critical	Warning	Exception	Actions
+ bcc_check_ntp	tcheck	0	3	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_network_top_connections_checker	tcheck	0	0	0	Details
+ bcc_disk_usage_checker	tcheck	0	0	0	Details
+ bcc_host_live_check	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among

them, the Critical, Warning, and Exception results are alerts. You need to pay special attention to them, especially the Critical and Warning results.

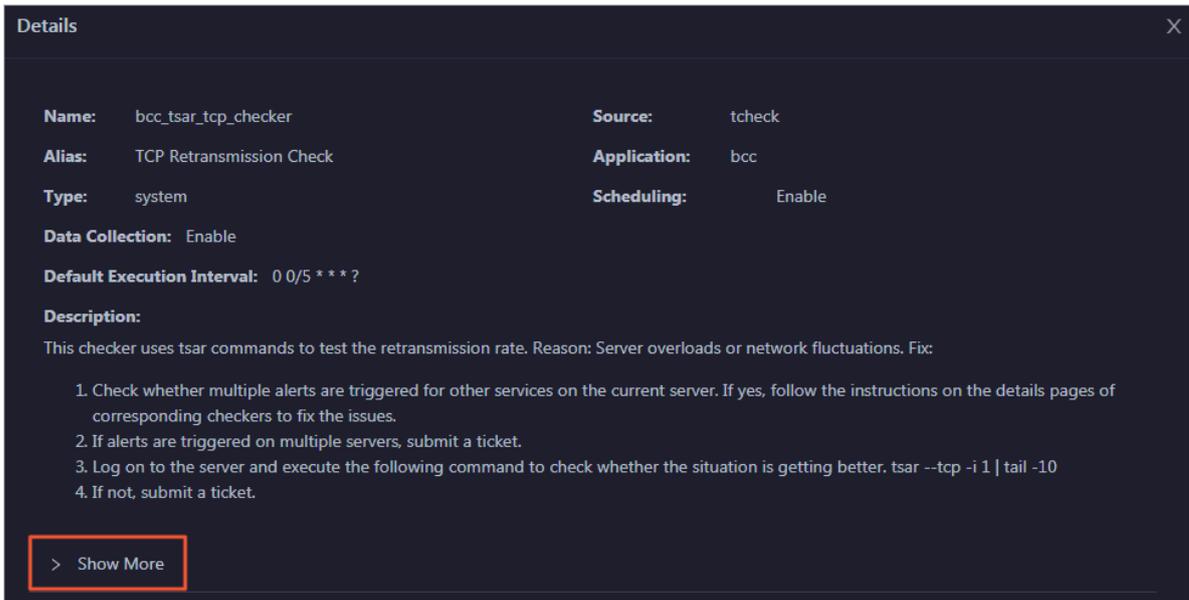
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

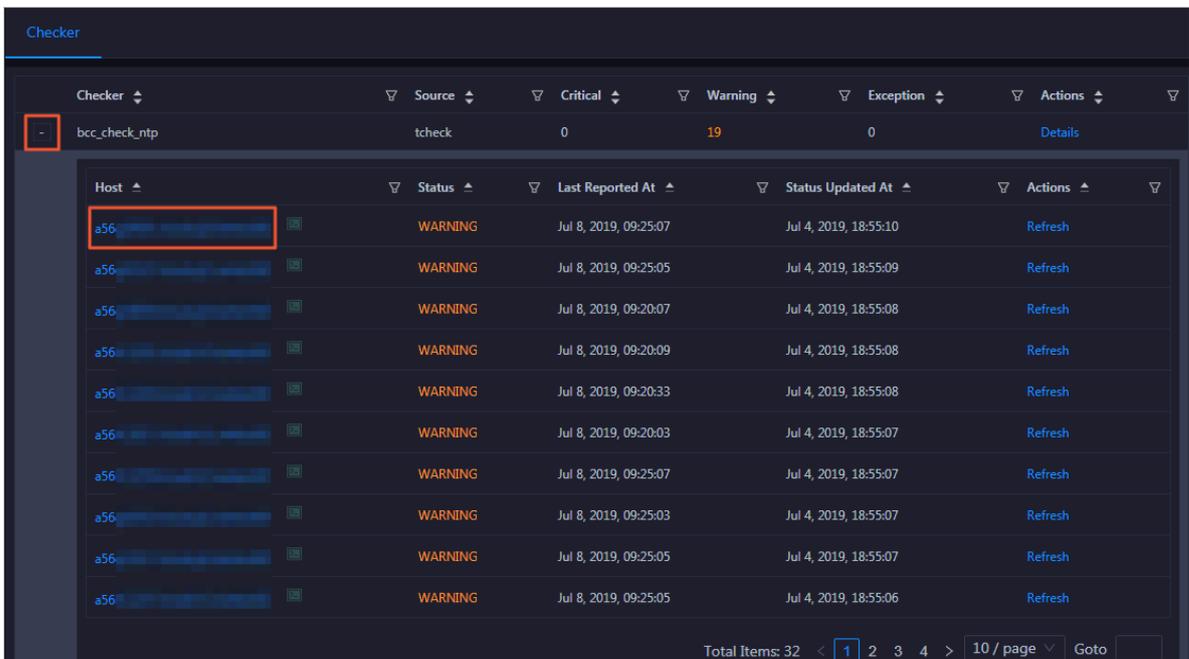


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

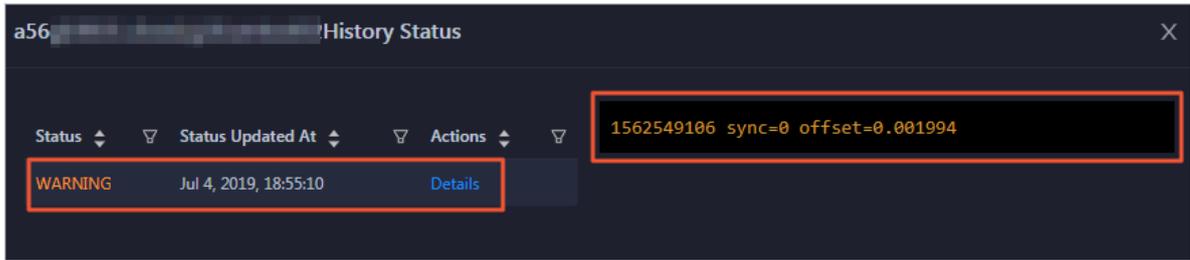
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.

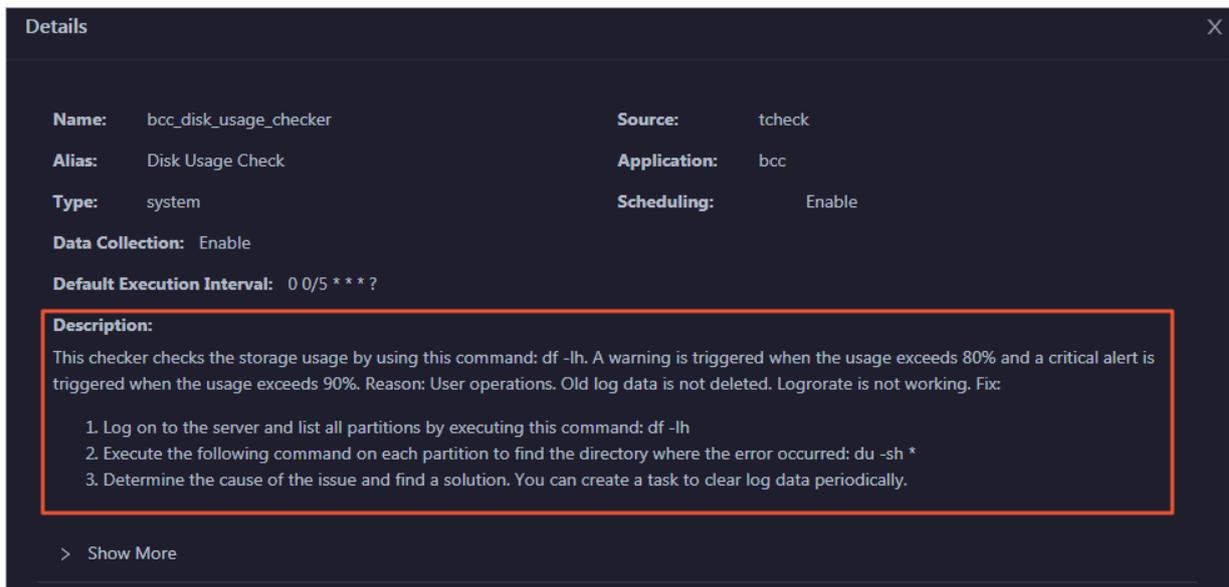


2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

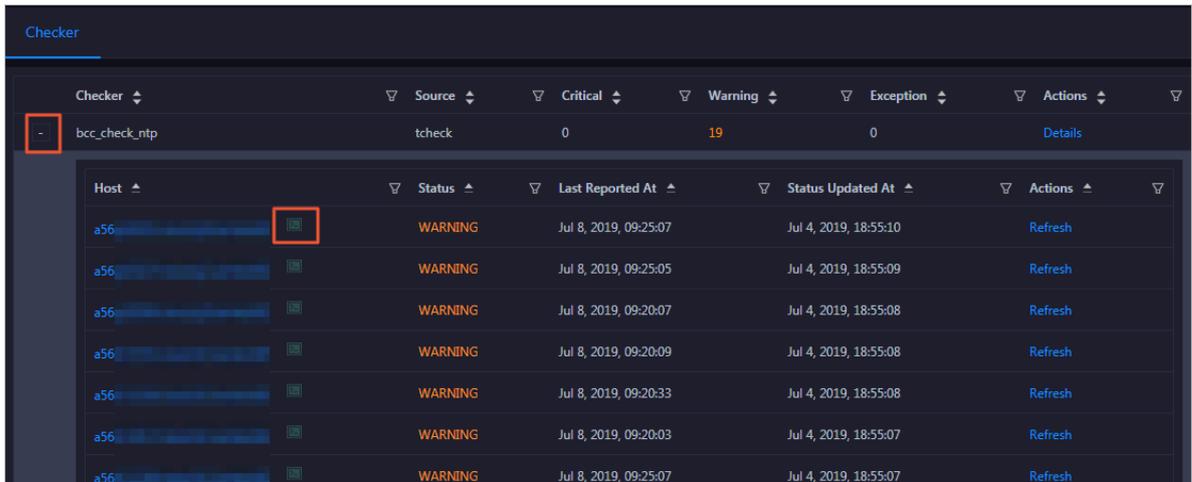
- On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



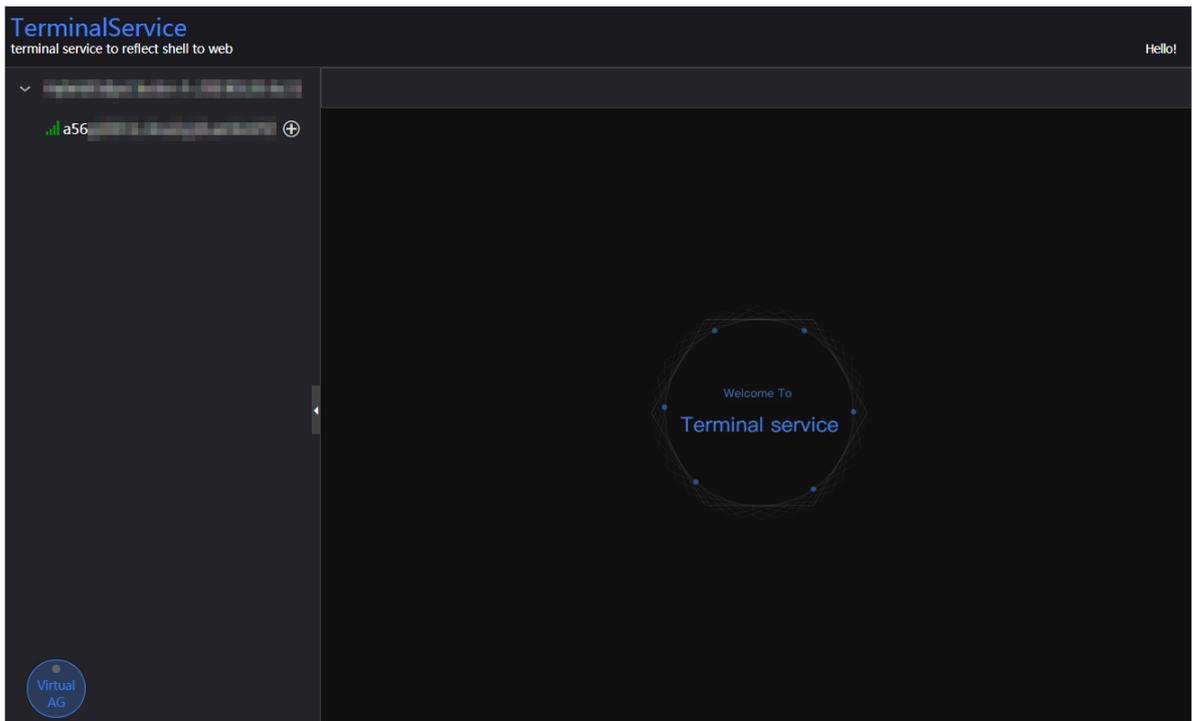
Log on to a host

- To log on to a host to clear alerts or perform other operations, follow these steps:

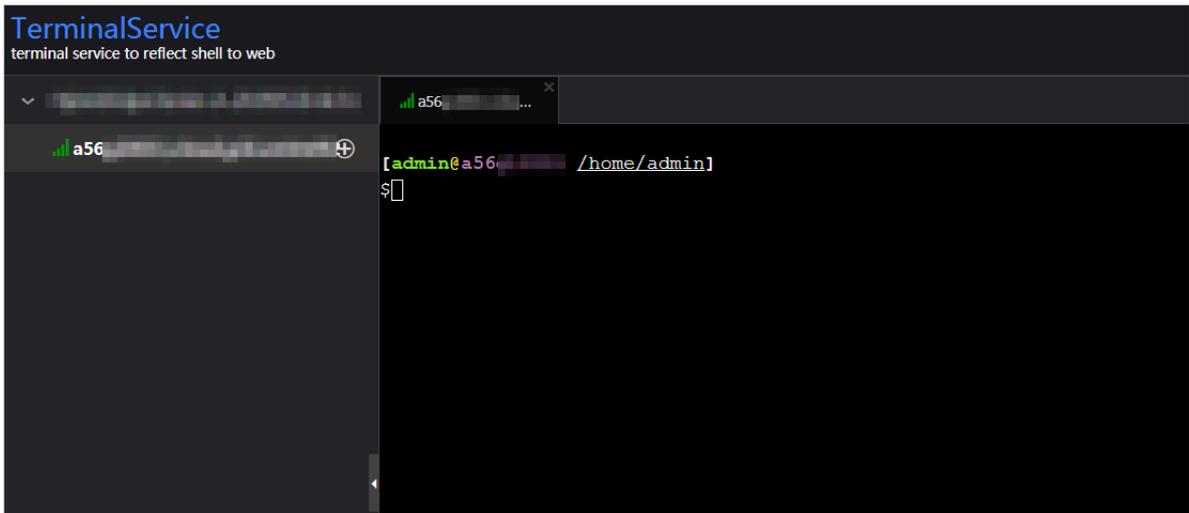
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon next to the name of a host with alerts. The TerminalService page appears.

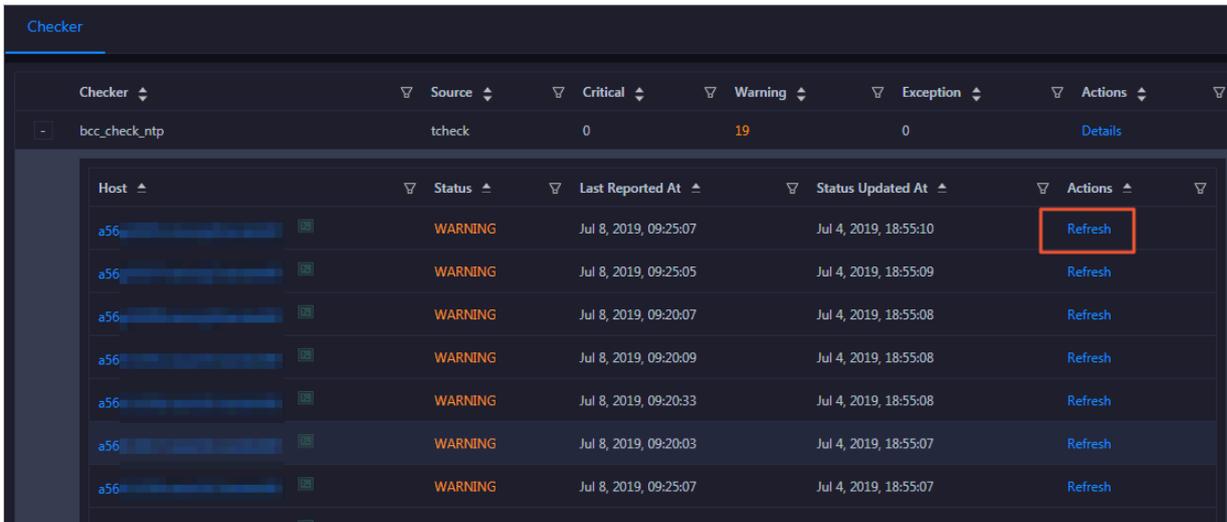


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



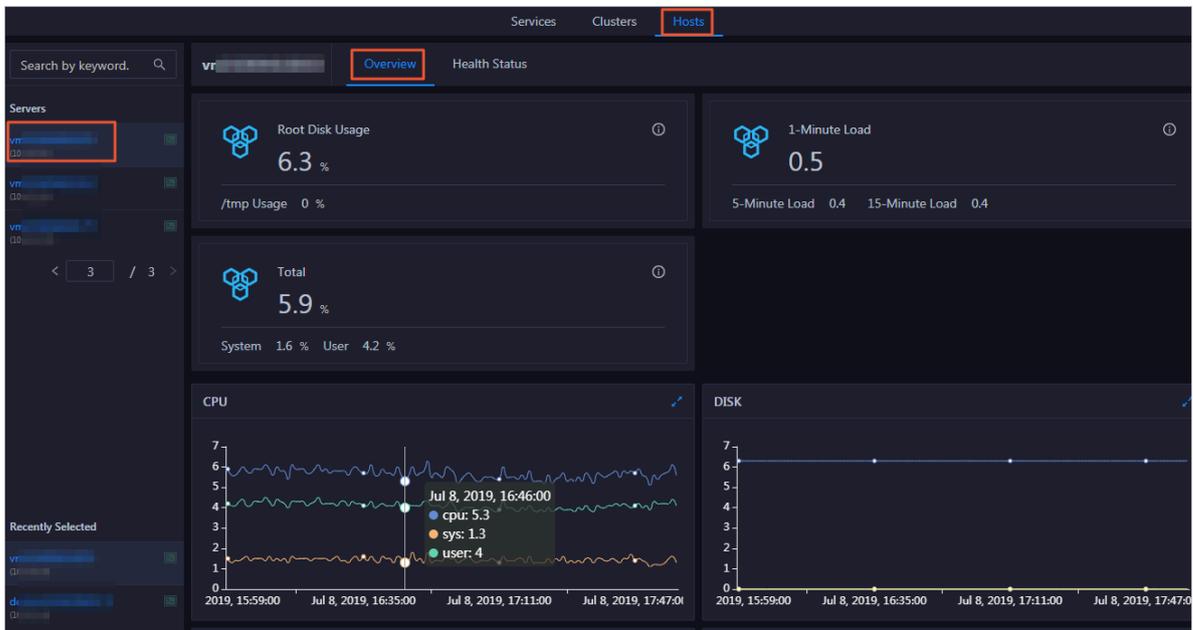
3.1.12.4 Host O&M

3.1.12.4.1 Host overview

The host overview page displays the overall running information about a host in a Machine Learning Platform for Artificial Intelligence (PAI) cluster. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Entry

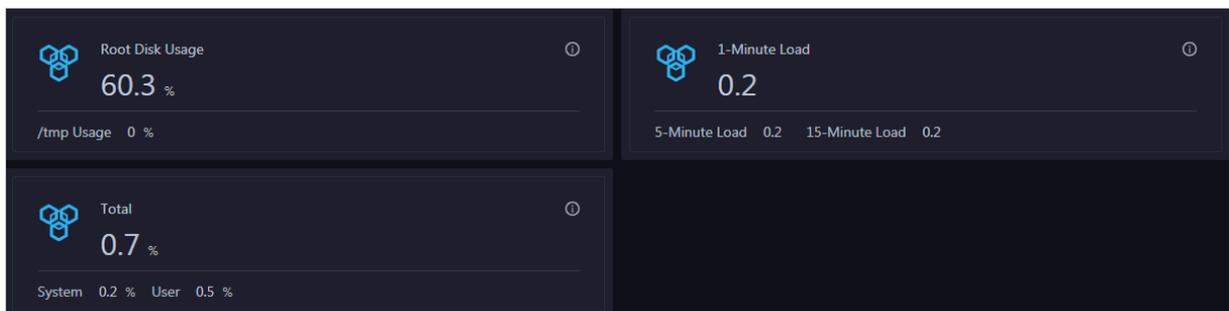
1. At the top of the O&M page, click the Hosts tab.
2. On the page that appears, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.



On the Overview page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.

Root Disk Usage, Total, and 1-Minute Load

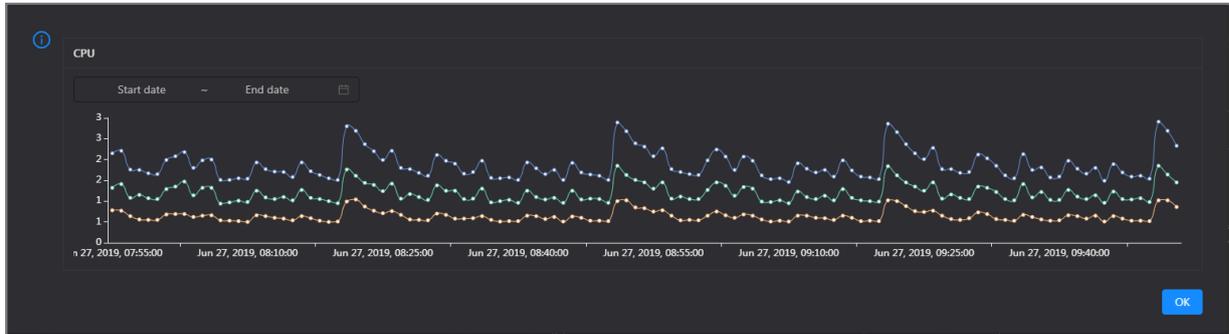
These sections display the root disk usage, total usage, and 1-minute load for the selected host. The Root Disk Usage section provides the usage of the `/tmp` directory. The Total section provides the system usage and user usage. The 1-Minute Load section provides the 1-minute, 5-minute, and 15-minute load averages.



CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

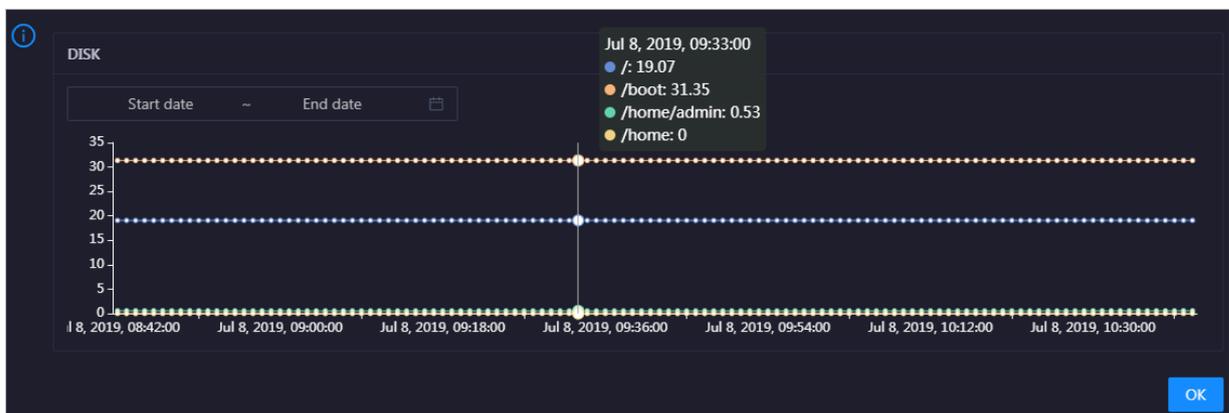


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the /, /boot, /home/admin, and /home directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



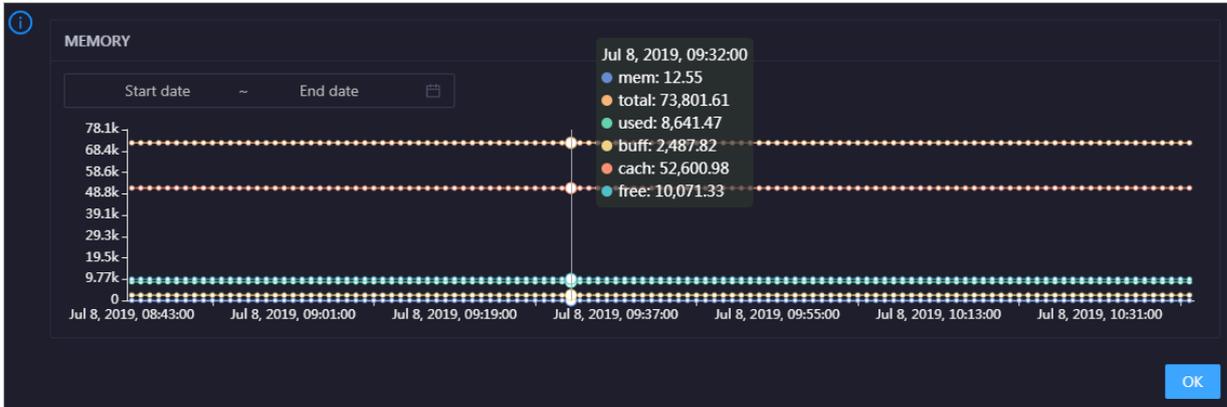
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size

of memory used by the page cache (cach), and available memory size (free) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

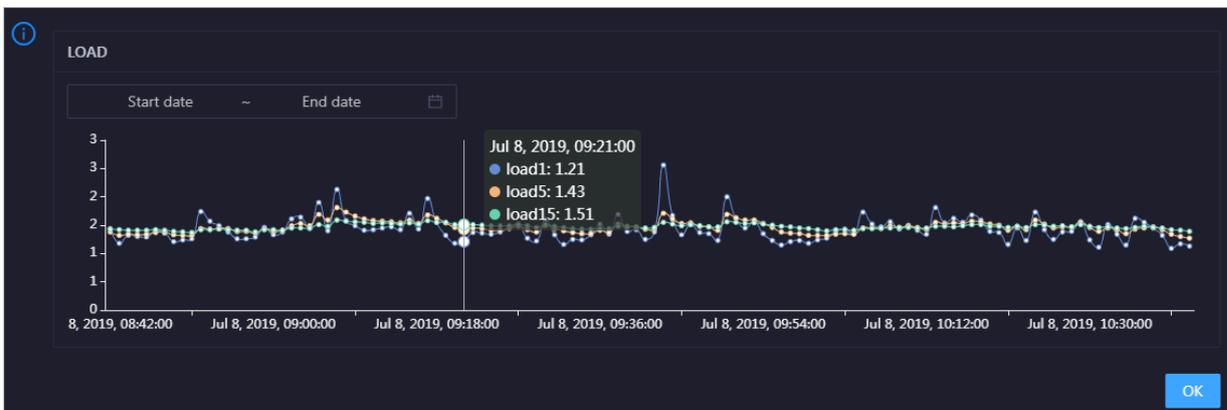


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



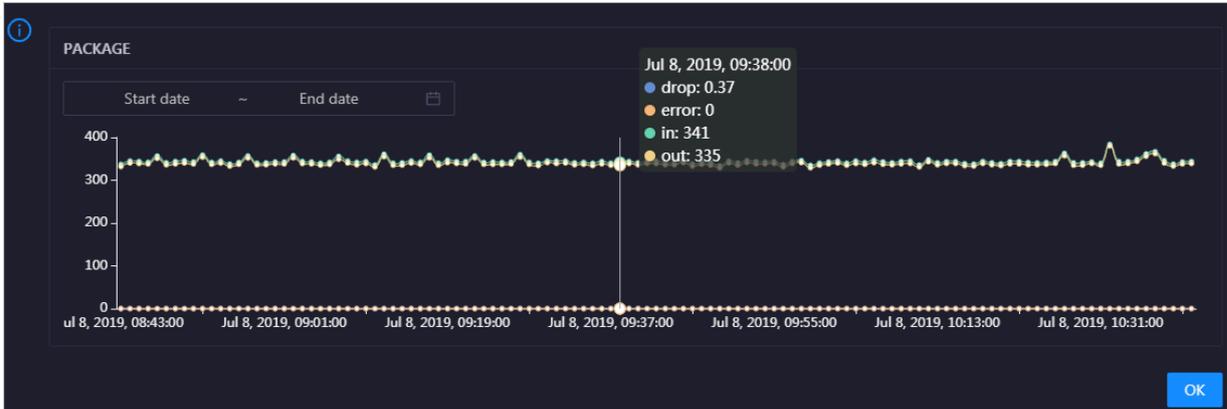
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for

the host over time in different colors. These trend lines reflect the data transmission status of the host.

Click  in the upper-right corner of the chart to zoom in it.

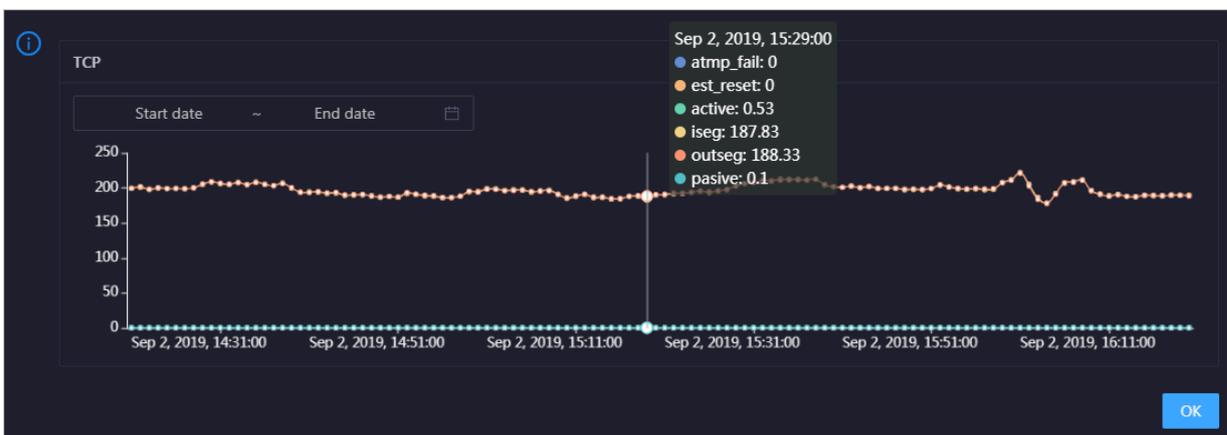


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the host over time in different colors. These trend lines reflect the TCP connection status of the host.

Click  in the upper-right corner of the chart to zoom in the chart.

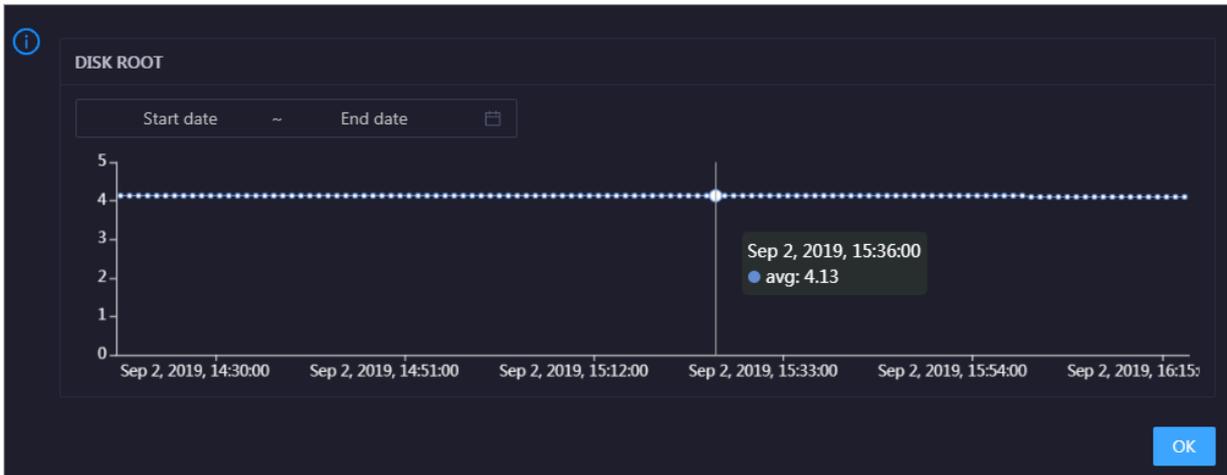


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the host in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the host over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the host in the specified period.

Health Check

This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.

Health Check [View Details](#)

Currently, 9 checkers are deployed on the service. 2 critical, 0 exception, and 0 warning alerts are reported.

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

Health Check History

This section displays a record of the health checks performed on the host.

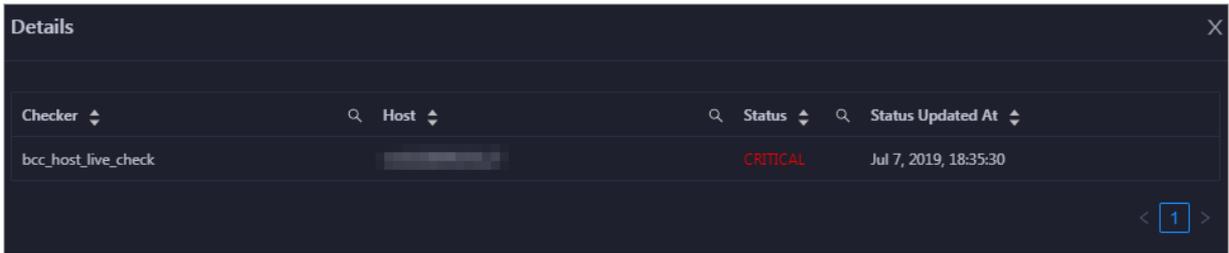
Health Check History [View Details](#)

Time	Event Content
Recently	1 alerts are reported by checkers.

< 1 >

Click [View Details](#) to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

You can click the event content of a check to view the exception items.

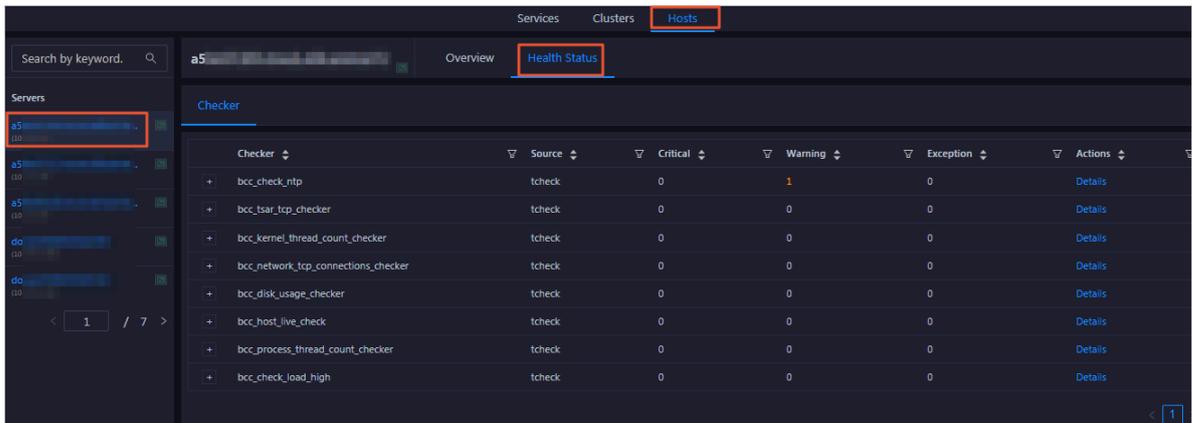


3.1.12.4.2 Host health

On the host health status page, you can view the checkers of the selected host, including the checker details, check results, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

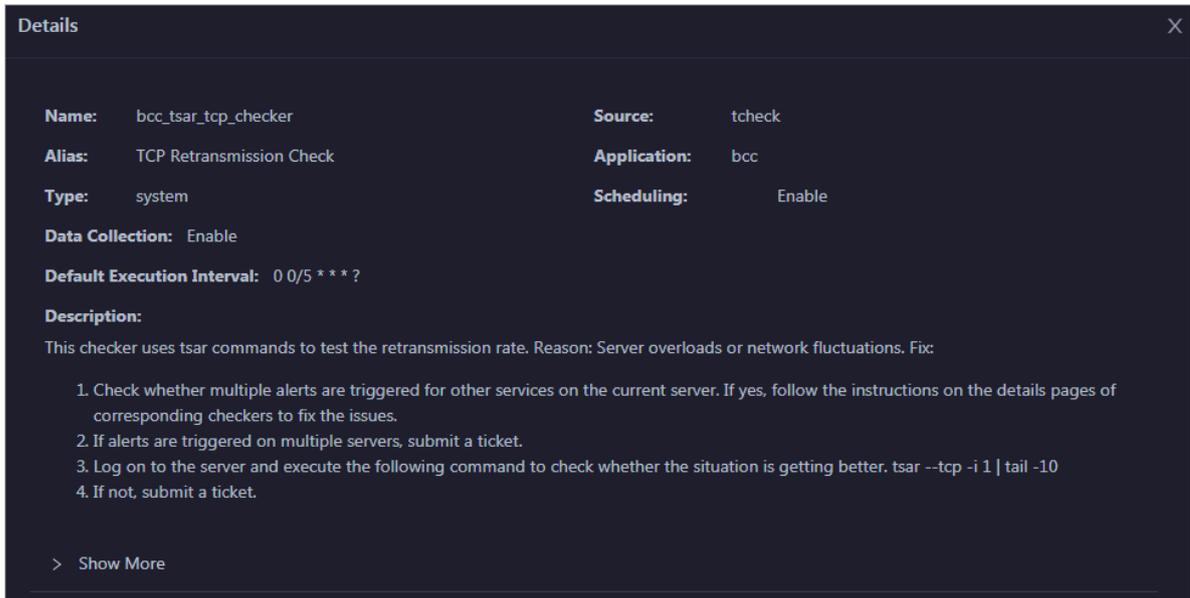
1. At the top of the O&M page, click the Hosts tab.
2. On the page that appears, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.



On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

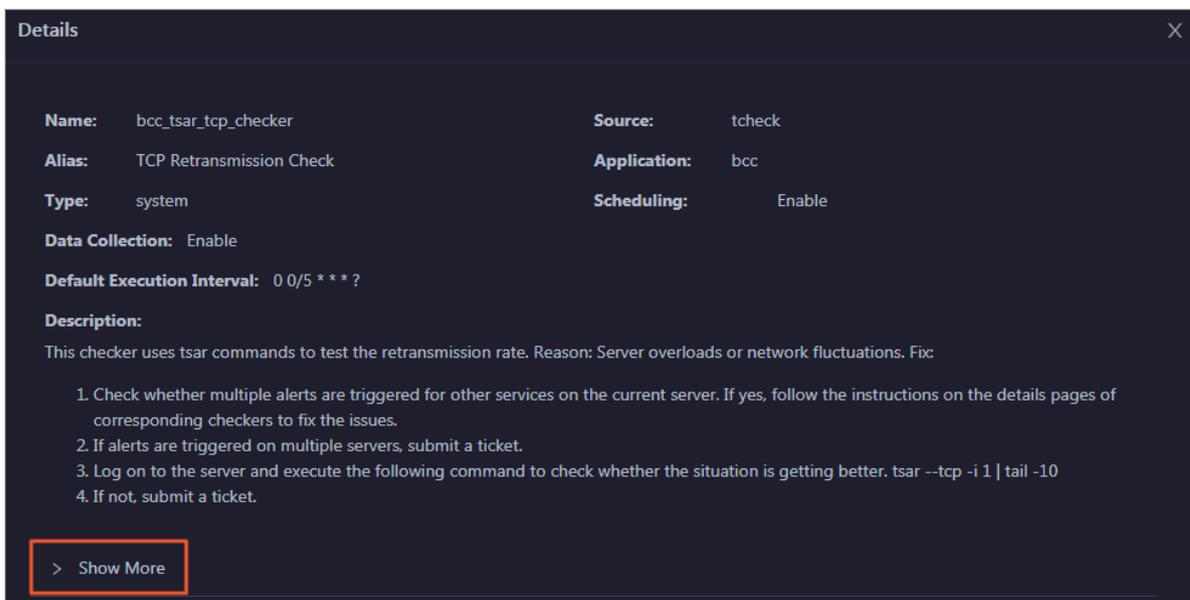
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

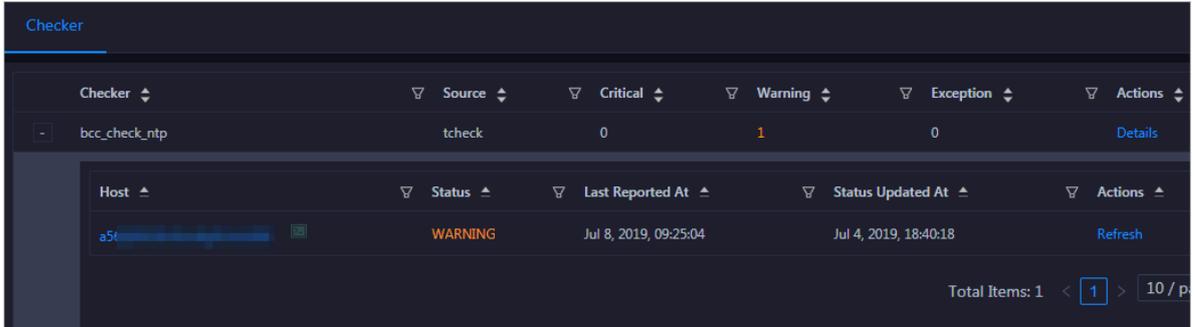


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

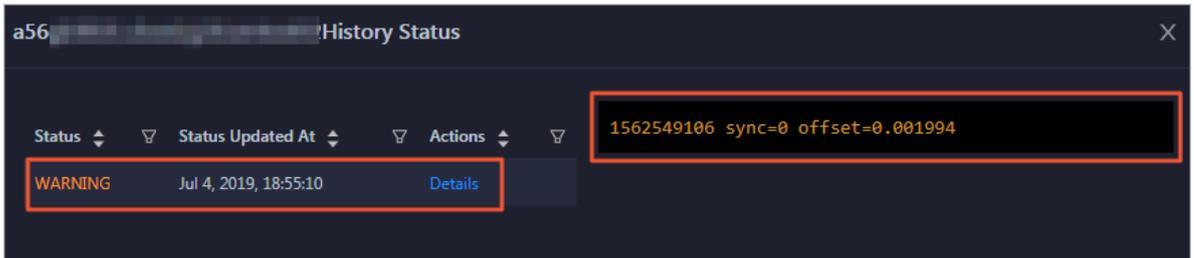
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

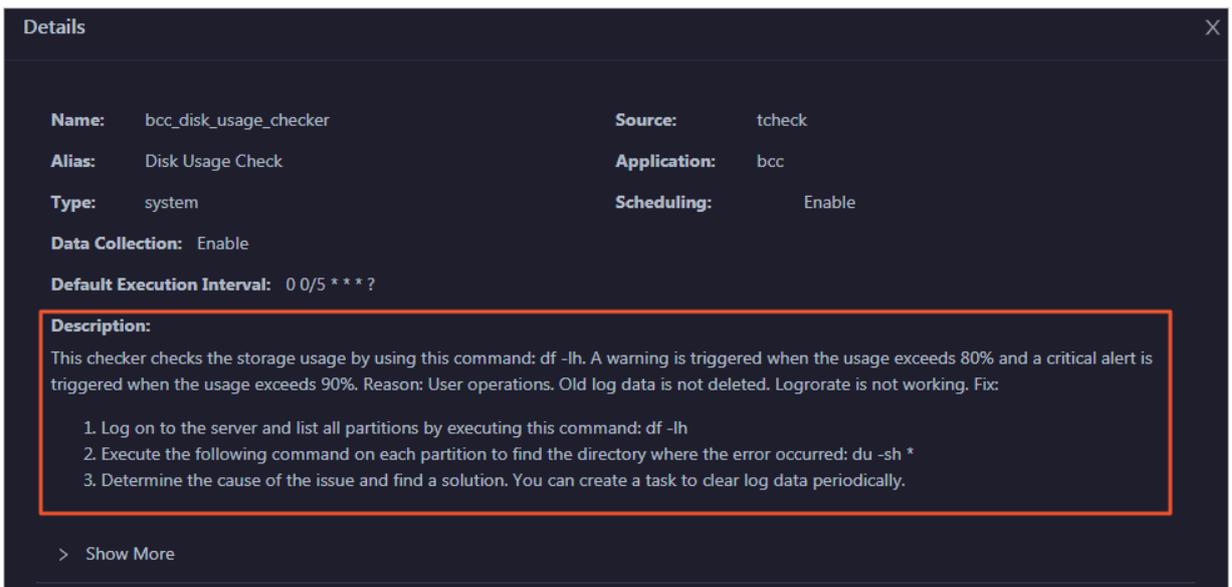


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

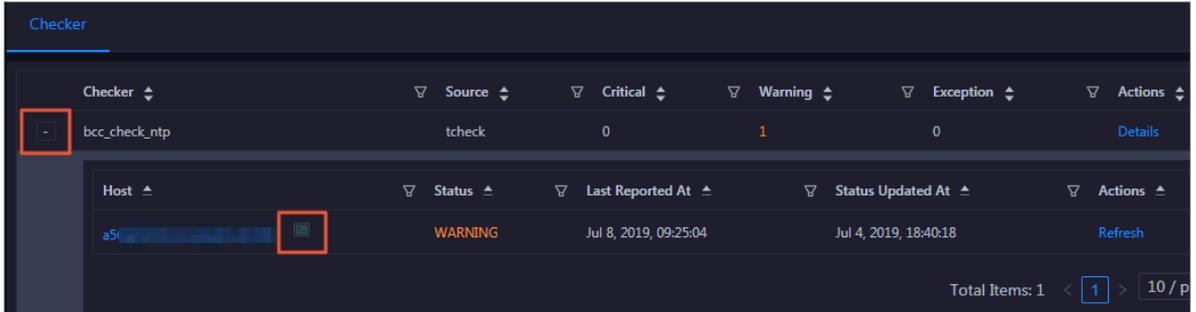
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



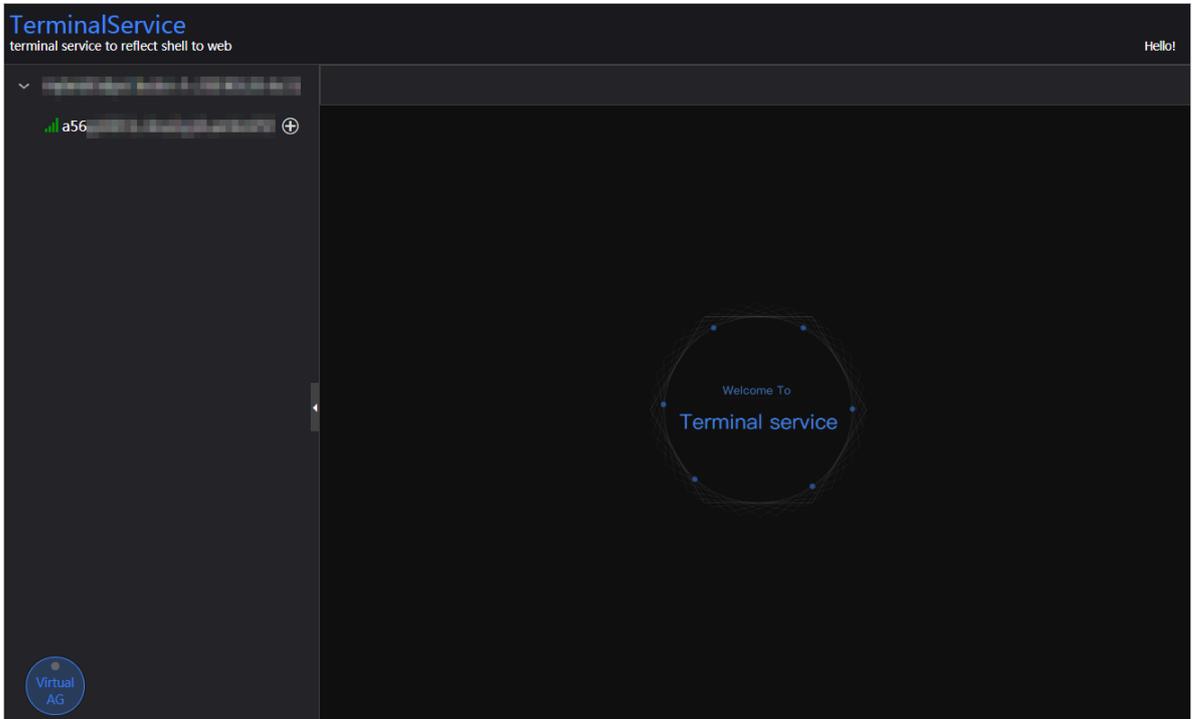
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

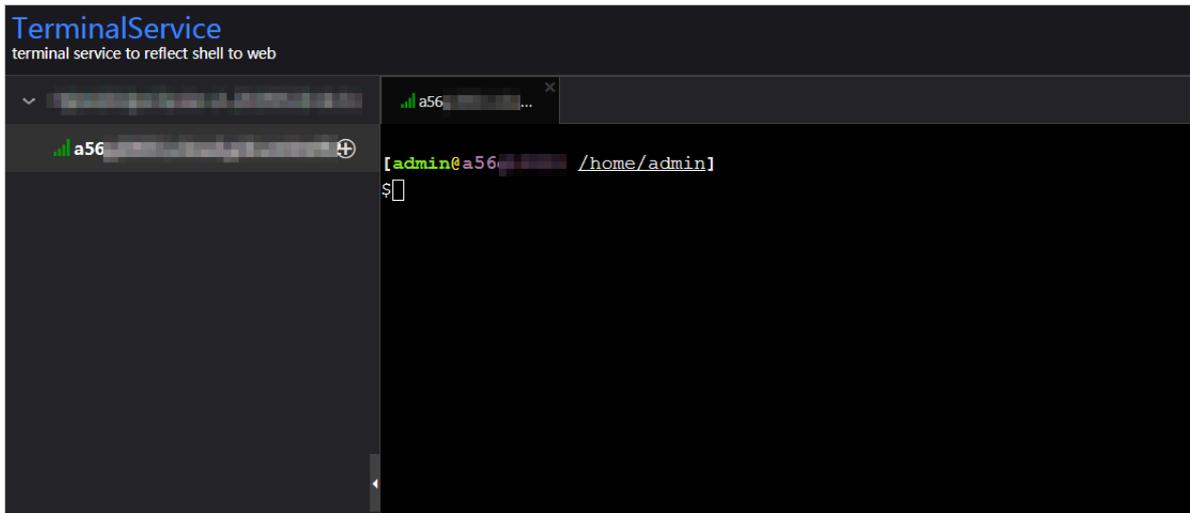
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

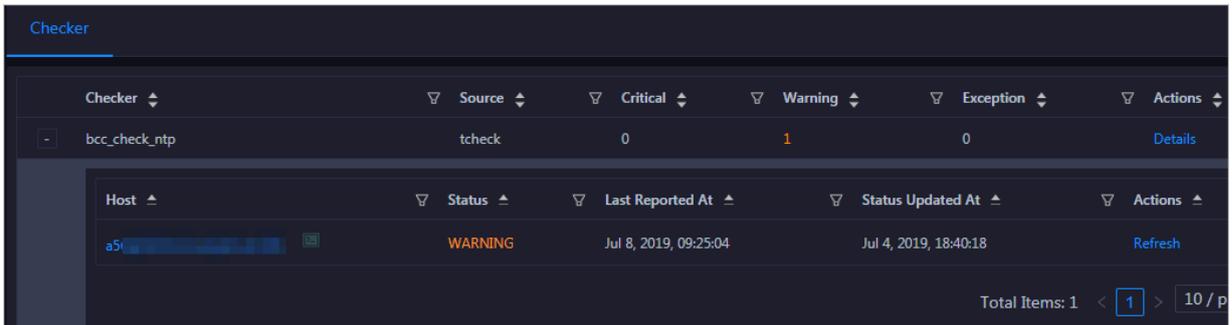


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.1.13 Management

3.1.13.1 Overview

The management module is the configuration and software management center of Apsara Bigdata Manager (ABM). It is an important functional module that supports and customizes O&M items for products.

The management module supports the following features:

- **Job execution and management:** You can generate jobs based on the scheme library to perform O&M operations on products.
- **Patch management:** You can deploy upgrade patches for various products.

- **Hot upgrade:** You can perform hot upgrades on the monitoring configuration and monitoring items of ABM so that services are not interrupted during the upgrade process.
- **Health management:** You can create health checkers and apply them to product hosts.
- **Operation audit:** You can view the records of job execution and other product O&M operations in ABM.

3.1.13.2 Jobs

3.1.13.2.1 Overview

This topic describes the job management interface and concepts related to jobs.

Apsara Bigdata Manager (ABM) runs jobs to perform O&M operations on big data products. Jobs, also known as product O&M tasks, are O&M operations performed on physical devices in the cluster. The job management interface consists of two pages: **Job Execution** and **Job Management**.

Concepts

Concepts related to jobs include:

- **Ordinary job:** jobs that can only be manually run.
- **Cron job:** jobs that are automatically run based on timer settings.
- **Scheme:** job templates provided by ABM. You can use schemes to generate jobs.
- **Atom:** step templates provided by ABM. You can use atoms as steps when generating jobs.
- **Ordinary step:** steps that you need to create when using schemes to generate jobs . Step types include the following: command execution, script execution, file push, API call, and manual step.
- **Atomic step:** steps that you can directly use when using schemes to generate jobs.

ABM provides common schemes and atoms that support most O&M scenarios.

Job Execution page

The screenshot shows the 'Job Execution' page with the following components:

- Scheme Library (Top 8):** A grid of eight job icons with labels: OdpsService_stop, OdpsService_start, sync_merge_data, ecs_gateway_scaleIn, Fuxi rm Readonly, Fuxi add Readonly, Pangu stop balance, and Pangu start balance.
- Cron Jobs (Top 8):** A grid of eight job icons with labels: odps_clean_history_data, odps_project_topn_tabl..., odps_smallfile_run, odps_smallfile_collect, odps_project_pangu_st..., odps_pangu_storage_in..., odps_cluster_res, and odps_collect_realtime_i...
- Ordinary Jobs:** A table with the following data:

Job Name	Created At	Modified At	Actions
OdpsService_stop	Jul 9, 2019, 15:49:28	Jul 9, 2019, 15:49:28	View Run History

The Job Execution page contains the following modules:

- **Ordinary Jobs:**

You can view and run ordinary jobs, and view their execution history.

You can search for a specific ordinary job.

- **Cron Jobs:**

You can enable, disable, view, and run cron jobs, and view their execution history.

You can search for a specific cron job.

- **Scheme Library (Top 8):** dynamically displays the top 8 most used schemes.

- **Cron Jobs (Top 8):** dynamically displays the top 8 most used cron jobs.

- **Execution History:**

You can view the execution history of ordinary and cron jobs.

You can search for the execution record of a specific job by multiple conditions.

Job Management page

Scheme Name	Created At	Modified At	Actions
OdpsService_stop	Apr 29, 2019, 16:52:14	Jun 5, 2019, 21:46:25	Run Generate Job History
OdpsService_start	Apr 29, 2019, 16:52:06	Jun 5, 2019, 21:46:13	Run Generate Job History
MaxCompute Chunkserver Scale-out	Apr 8, 2019, 16:41:45	May 27, 2019, 21:50:43	Run Generate Job History
MaxCompute Chunkserver Scale-in	Apr 8, 2019, 16:41:41	May 27, 2019, 21:50:36	Run Generate Job History
DataWorks Gateway Scale-out	Apr 8, 2019, 16:36:59	May 27, 2019, 21:50:28	Run Generate Job History
Dataworks Gateway Scale-in	Apr 8, 2019, 16:36:51	May 27, 2019, 21:50:16	Run Generate Job History
Change Bcc Dns-Vip Relation For Disaster Recovery	Apr 8, 2019, 16:36:21	May 21, 2019, 19:29:27	Run Generate Job History
ODPS_Stop_Service_Mode	Apr 8, 2019, 16:57:02	Apr 12, 2019, 16:05:37	Run Generate Job History
ODPS_Start_Service_Mode	Apr 8, 2019, 16:43:38	Apr 12, 2019, 15:27:02	Run Generate Job History
sync_merge_data	Apr 8, 2019, 16:45:13	Apr 8, 2019, 16:45:13	Run Generate Job History

The Job Management page provides the following features:

- You can generate and run jobs based on schemes and view the execution history of schemes.
- You can search for a specific scheme.
- You can view schemes in grid or list mode.

3.1.13.2.2 Jobs

3.1.13.2.2.1 Run a job from a scheme

When you perform O&M operations, you can directly run jobs from schemes that meet your requirements. This enables you to quickly perform product O&M jobs.

Prerequisites

You must have an ABM administrator account.

Context

When you run a job from a scheme, you need to specify the Target Group and Global Variable parameters. The other parameters cannot be modified. If you want to modify the parameters, see [Create a job from a scheme](#).

Running a job from a scheme is a one-time operation and does not generate a job on the Ordinary Jobs tab. You can view the history operations on the Execution History tab. For more information, see [View the execution history](#).

Procedure

1. [Log on to the ABM console](#).
2. Click Management in the upper-right corner. On the page that appears, click Jobs in the left-side navigation pane.
3. Run a job by using one of the following methods:
 - In the Scheme Library (Top 8) section, select a scheme.



Note:

This method only allows you to choose a scheme from the top 8 most frequently used schemes.

- On the Jobs page, click the Job Management tab, and then click Run in the Actions column of a scheme in the Schemes list.

4. On the Run from Scheme page, you need to set Target Group and Variable Name as needed.

The instructions for setting Target Group and Variable Name are shown in [Table 3-1: Job parameters](#).

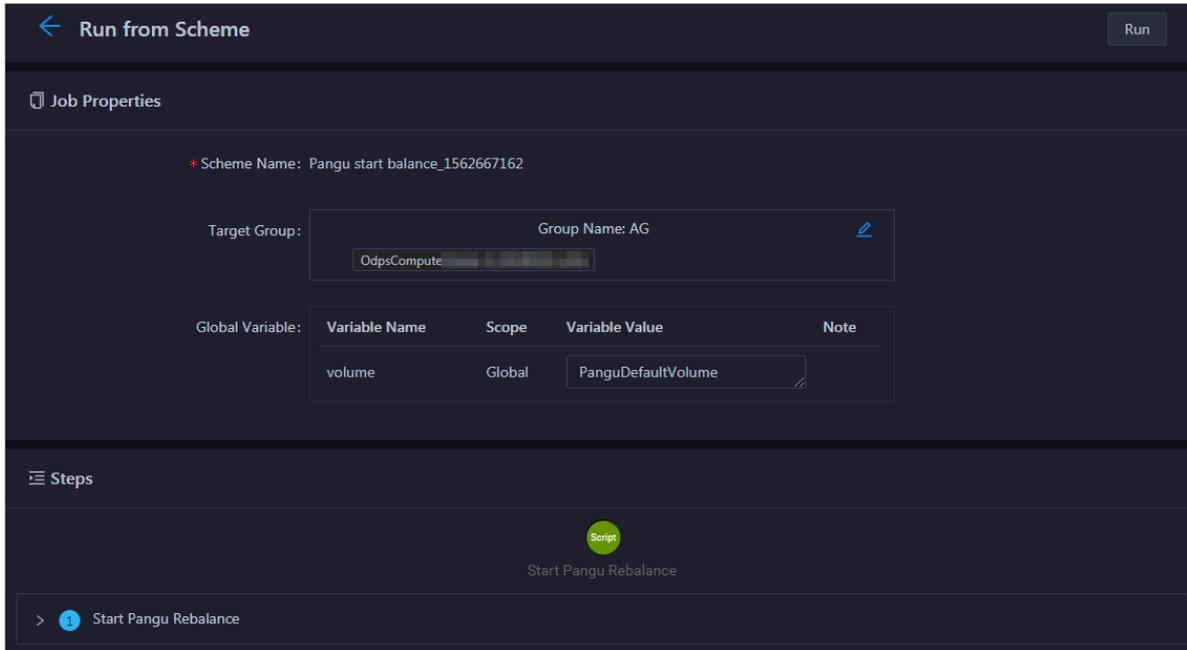
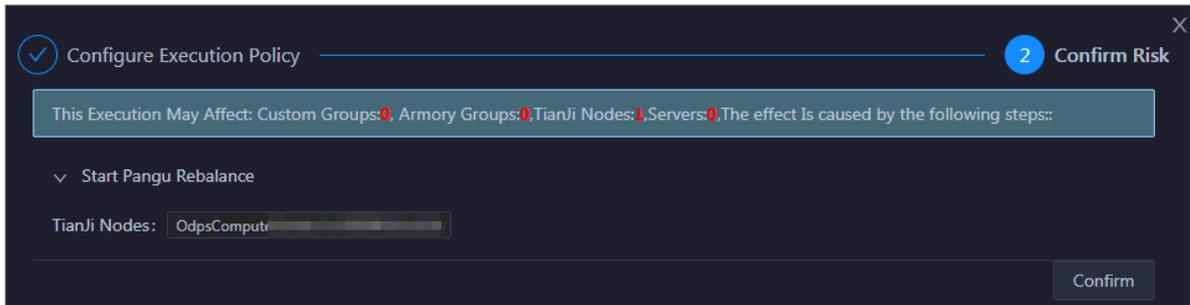


Table 3-1: Job parameters

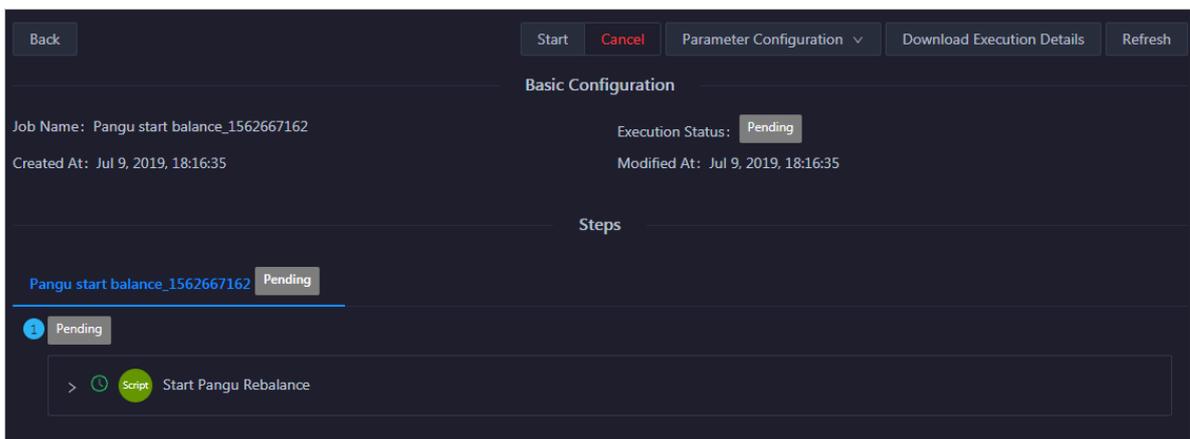
Parameter	Description
Target Group	<p>A collection of target nodes on which the operations are performed. After you have added nodes to target groups, you can select a value for Target Group based on your needs when you configure the steps.</p> <p>Click  next to the target group, and set the nodes to be included in the target group as needed. When you add a node, you can either select the name of the node in Apsara Infrastructure Management Framework or enter the IP address of the node under Servers.</p>
Global Variable	<p>If global variables are set in the scheme, you need to enter the variable value.</p>

5. After you have configured the preceding parameters, click **Run** in the upper-right corner.
6. Confirm the job risks in the displayed dialog box, and click **Confirm Execution**.



After you have confirmed, a record is automatically generated on the **Execution History** page. For more information, see [View the execution history](#).

7. On the job execution page, click **Start** at the top to start the execution.



If you do not perform any operation and exit the job execution page, you can find a job record on the **Execution History** page. Click **View** to go to the job execution page again.

3.1.13.2.2.2 Create a job from a scheme

This topic describes how to generate a job from a scheme. You can generate both ordinary and cron jobs from schemes.

Prerequisites

You must have an Apsara Bigdata Manager (ABM) administrator account.

Context

ABM allows you to create both ordinary and cron jobs from schemes. Settings for creating an ordinary job and a cron job are similar, but a schedule must be created for a cron job.

Procedure

1. [Log on to the ABM console.](#)
2. Click **Management** in the upper-right corner. On the page that appears, click **Jobs** in the left-side navigation pane.
3. On the **Jobs** page, click the **Job Management** tab, and click **Generate Job** in the **Actions** column of a scheme in the **Schemes** list.
4. On the **Create Job** page, set the parameters in the **Job Properties** and **Steps** sections as needed.

For more information about the parameter configuration of **Job Properties**, see [Table 3-2: Description of job properties.](#)

The screenshot shows the 'Create Job' interface. At the top, there is a 'Save' button. Below it, the 'Job Properties' section is visible. It includes a 'Job Type' selector with 'Ordinary Job' and 'Cron Jobs' options. The 'Job Name' field is filled with 'Pangu start balance'. The 'Target Group' section shows a 'Group Name: AG' with edit and delete icons, and a '+ Create Group' button. The 'Global Variable' section is a table with columns for Variable Name, Scope, Variable Value, Note, and Actions. One variable is listed: 'volume' with a 'Global' scope and 'PanguDefaultVolume' value. There is an '+ Add Variable' button at the bottom.

Table 3-2: Description of job properties

Parameter	Description
Job Type	<p>The type of the job.</p> <ul style="list-style-type: none"> • Ordinary Job: jobs that must be manually run. • Cron Jobs: jobs that automatically run based on a schedule. You can enter a cron expression or click Configure Cron Job to create a schedule. <p>Cron expressions are based on crontab commands. If you are new to crontab commands, click Configure Cron Job to quickly set up a schedule.</p>

Parameter	Description
Job Name	The name of the job. Set the job name based on the functionality of the job to be created so that the user understands what it is and can search for it.
Target Group	<p>A collection of target nodes on which the operations are performed. After you have added nodes to the target groups, you can select the target group based on your needs when you configure the steps.</p> <p>After you have created a group, click  to add nodes to the group. When you add a node, you can either select the name of the node in Apsara Infrastructure Management Framework or enter the IP address of the node under Servers.</p>
Global Variable	Click Add Variable and set the parameters in the dialog box that appears. The Scope parameter is used to set the scope of the variable. If it is set Global, it is valid for the entire job. If you select a certain step, it is only valid for this step.

5. On the Create Job page, add steps as needed.



The steps include ordinary steps and atomic steps.

- **Atomic Step:** a range of built-in steps provided by the system.
- **Ordinary Step:** Ordinary steps include multiple types. Choose the required type and set the parameters accordingly. [Table 3-3: Parameters of command execution steps](#), [Table 3-4: Parameters of script execution steps](#), [Table 3-5: Parameters of file push steps](#),

Table 3-6: Parameters of API call steps, and *Table 3-7: Parameters of manual steps* describe the parameters for different types of ordinary steps.

Table 3-3: Parameters of command execution steps

Section	Parameter	Description
N/A	Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Basic configuration	Target Node Group	The group of nodes on which the step is performed.
	Commands	The commands to be executed in this step.
	User Identity	The user that executes this step on the nodes, with a default setting of admin.
	Description	The description of the step.
Advanced configuration	Input Context	Enable the input context to obtain the output of the previous step. When enabled, this step reads the file specified by the <i>\$contextInput</i> variable to obtain the context.
	Output Context	Enable this option if you need to export the context to the next step. When enabled, this step writes the context to the file specified by the <i>\$contextOutput</i> variable to export the context.
	Timeout Period	The maximum time allowed to execute the step. If the step is not complete before the time runs out, the execution is stopped and you are notified that the operation is timed out. The default value is 60 seconds.
	Retries	The number of times to retry a target after a failure or timeout. The default value is 0.

Section	Parameter	Description
	Retry Interval	<p>The interval between two executions. The default value is 300 seconds.</p> <p>The retry interval is the period of time between the last timeout (or failure) and the next try.</p>

Table 3-4: Parameters of script execution steps

Section	Parameter	Description
N/A	Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Basic configuration	Target Node Group	The group of nodes on which the step is performed.
	Script Content	<p>Write the script based on the actual O&M requirements. Currently, Shell and Python are supported.</p> <p>You can write new scripts or upload local scripts to configure the script content.</p>
	User Identity	The user that executes this step on the nodes, with a default setting of admin.
	Description	The description of the step.
Advanced configuration	Input Context	Enable the input context to obtain the output of the previous step. When enabled, this step reads the file specified by the <code>\$contextInput</code> variable to obtain the context.
	Output Context	Enable this option if you need to export the context to the next step. When enabled, this step writes the context to the file specified by the <code>\$contextOutput</code> variable to export the context.

Section	Parameter	Description
	Timeout Period	<p>The maximum time allowed to execute the step. If the step is not complete before the time runs out, the execution is stopped and you are notified that the operation is timed out.</p> <p>The default value is 60 seconds.</p>
	Retries	<p>The number of times to retry a target after a failure or timeout.</p> <p>The default value is 0.</p>
	Retry Interval	<p>The interval between two executions. The default value is 300 seconds.</p> <p>The retry interval is the period of time between the last timeout (or failure) and the next try.</p>

Table 3-5: Parameters of file push steps

Parameter	Description
Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Target Node Group	The group of nodes to which the file is pushed.
Target Path	The directory to which the file is pushed.
File Permission	The permission of the file.
File Owner	The owner of the file.

Parameter	Description
File Content	<p>Enter the file content in the Edit interface or upload a local file.</p> <p>After you have entered or uploaded the content, click Edit to go to File Management and specify the file name.</p>

Table 3-6: Parameters of API call steps

Parameter	Description
Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Target URL	The URL of the API.
HTTP Method	<p>The type of request that you want to send.</p> <ul style="list-style-type: none"> • GET: Query. • POST: Create. • PUT: Modify. • DELETE: Delete.
Content Format	The Content-Type field of the header in the HTTP packet. Select from the drop-down list.
APP NAME	APP NAME and APP KEY are included in the request to call APIs for authenticating permissions.
APP KEY	
BODY	The body of the HTTP request.
Timeout Period	<p>The maximum time allowed to execute the step. If the step is not complete before the time runs out, the execution is stopped and you are notified that the operation is timed out.</p> <p>The default value is 60 seconds.</p>
Retries	<p>The number of times to retry a target after a failure or timeout.</p> <p>The default value is 0.</p>

Parameter	Description
Retry Interval	<p>The interval between two executions. The default value is 300 seconds.</p> <p>The retry interval is the period of time between the last timeout (or failure) and the next try.</p>

Table 3-7: Parameters of manual steps

Parameter	Description
Step Name	The name of the step. Enter a step name that reflects the functionality of the step.
Document Content	The instructions to help relevant engineers complete this step.

6. To change the order of steps, click Sort in the upper-right corner of the Steps list and drag the steps to put them into the correct order.
7. After you have set the preceding parameters, click Save in the upper-right corner.

Result

If you created an ordinary job, it is displayed in the Ordinary Jobs list. If you created a cron job, it is displayed in the Cron Jobs list.

What's next

- If you created an ordinary job, you need to run it manually. For more information, see [Manually run a job](#).
- If you created a cron job, you need to enable it. For more information, see [Enable or disable a cron job](#). You can also manually run a cron job. For more information, see [Manually run a job](#).

3.1.13.2.2.3 Enable or disable a cron job

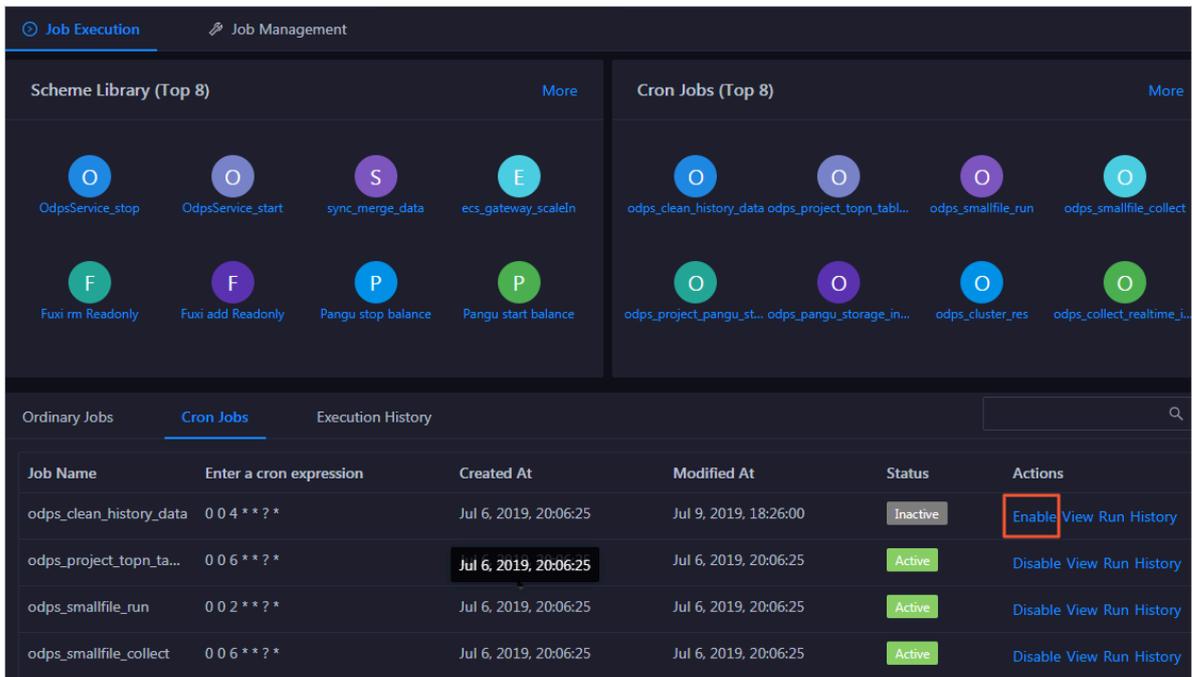
When a cron job is generated from a scheme, the job is disabled by default. You must manually enable it. If you do not need the cron job to run during a specified time period, you can manually disable it.

Prerequisites

You must have an ABM administrator account.

Procedure

1. [Log on to the ABM console.](#)
2. Click Management in the upper-right corner. On the page that appears, click Jobs in the left-side navigation pane.
3. On the Job Execution page, click Cron Jobs.



4. On the Cron Jobs page, you can enable or disable a cron job.

- To enable a cron job in the inactive status, click Enable in the Actions column of the cron job.

After a cron job is enabled, its status changes to Active. The Enable button is replaced by Disable.

- To disable a cron job in the active status, click Disable in the Actions column of the cron job.

After a cron job is disabled, its status changes to Inactive. The Disable button is replaced by Enable.

3.1.13.2.2.4 Manually run a job

After you have created an ordinary job, you must manually run the job in order to perform O&M operations on the product. You can also manually run a cron job.

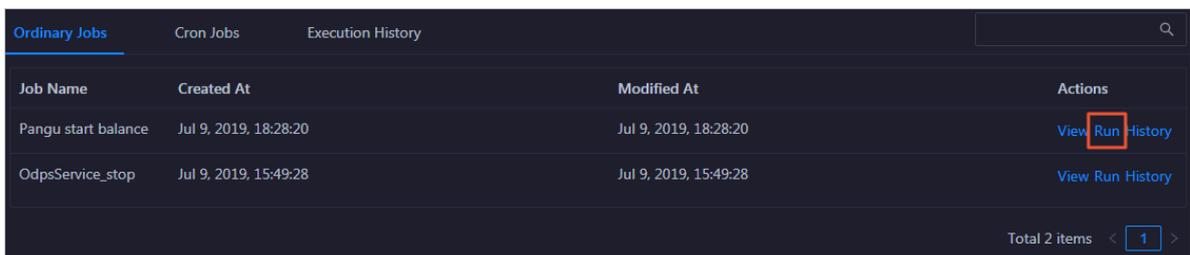
Prerequisites

You must have an ABM administrator account.

Procedure

1. [Log on to the ABM console](#).
2. Click **Management** in the upper-right corner. On the page that appears, click **Jobs** in the left-side navigation pane.
3. Click **Ordinary Jobs** on the **Job Execution** page.

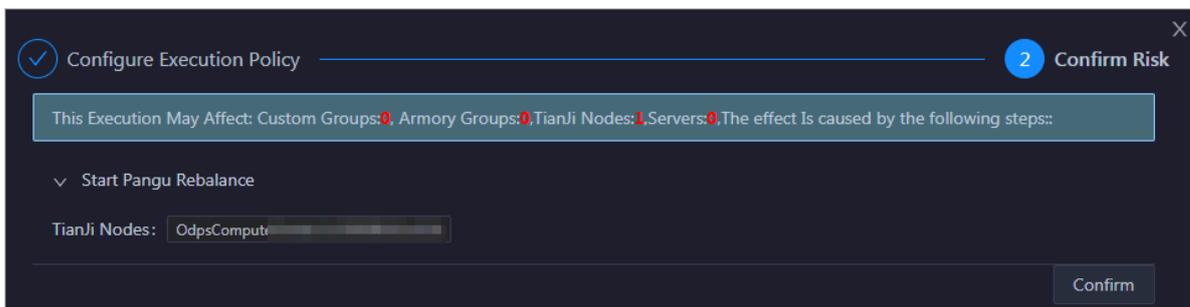
If you need to manually run a cron job, click **Cron Jobs**. The procedure to manually run a cron job is the same as that of an ordinary job. This topic takes ordinary jobs as an example.



Job Name	Created At	Modified At	Actions
Pangu start balance	Jul 9, 2019, 18:28:20	Jul 9, 2019, 18:28:20	View Run History
OdpsService_stop	Jul 9, 2019, 15:49:28	Jul 9, 2019, 15:49:28	View Run History

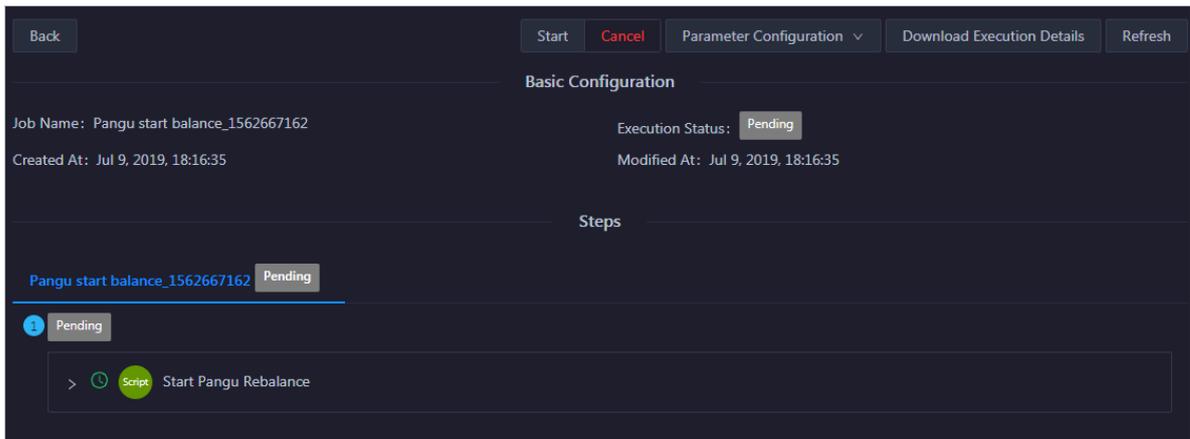
Total 2 items < 1 >

4. In the Ordinary Jobs list, click **Run** in the Actions column of a job.
5. Confirm the job risks in the dialog box that appears, and click **Confirm**.



After you have confirmed, a record is automatically generated on the **Execution History** page. For more information, see [View the execution history](#).

6. On the job execution page, click **Start** at the top to start the execution.



You can find the record about a job on the Execution History page, and click **View** to go to the detailed execution page.

3.1.13.2.2.5 View jobs

After you have created an ordinary job or a cron job, you can view job details, save the job as a scheme, and run the job in the jobs list.

Prerequisites

You must have an ABM administrator account.

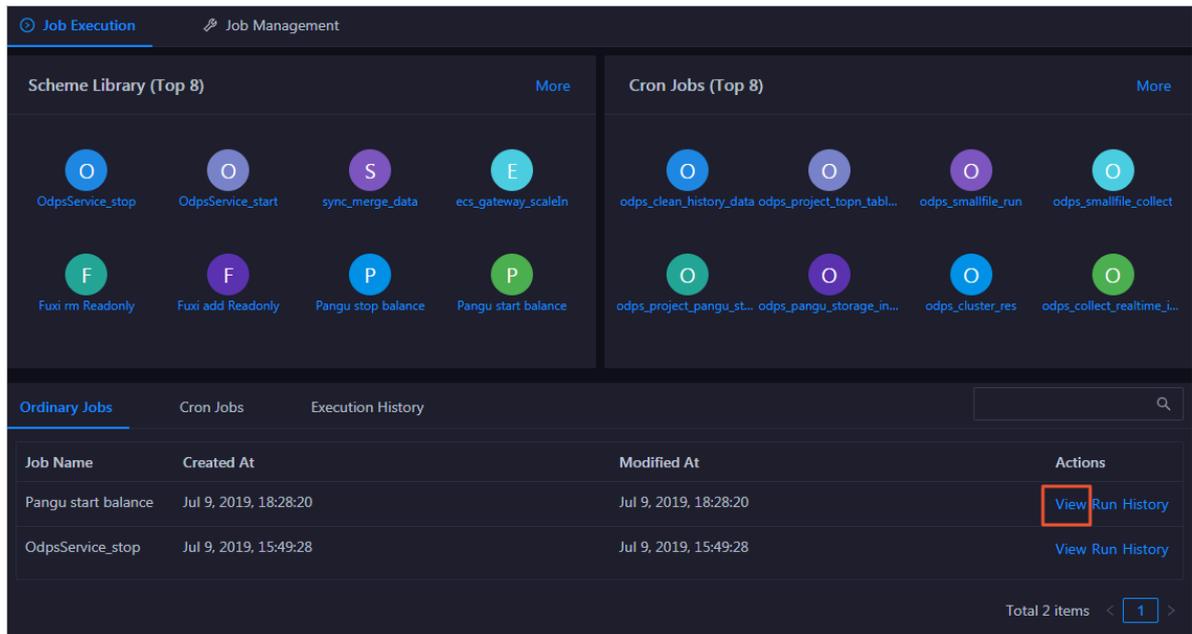
Context

The topic describes how to view ordinary jobs. You can follow the same procedure to view cron jobs.

Procedure

1. [Log on to the ABM console.](#)
2. Click **Management** in the upper-right corner. On the page that appears, click **Jobs** in the left-side navigation pane.
3. Click **Ordinary Jobs** on the Job Execution page.

4. Click View in the Actions column of an ordinary job to view its job details.



Job Name	Created At	Modified At	Actions
Pangu start balance	Jul 9, 2019, 18:28:20	Jul 9, 2019, 18:28:20	View Run History
OdpsService_stop	Jul 9, 2019, 15:49:28	Jul 9, 2019, 15:49:28	View Run History

3.1.13.2.2.6 View the execution history of a job

Apsara Bigdata Manager (ABM) allows you to view the execution history of a specific job to learn the execution status of it.

Prerequisites

You must have an ABM administrator account.

Context

After you confirm to run a job, ABM generates log information for the job execution. You can learn the execution status by using the log data.

The Execution History page provides the following features:

- Provides information such as the trigger mode, current status, start time, and end time of the job.
- Provides job execution details and parameter setting information, and allows you to download execution details.
- Allows you to perform different operations based on job status. For example, you can perform the Start operation on a job in the Not Started status and the Retry operation on a job in the Exception status.

This topic provides information on how to view the execution history of an ordinary job. You can follow a similar procedure to view the execution history of cron jobs.

Procedure

1. [Log on to the ABM console](#).
2. Click Management in the upper-right corner. On the page that appears, click Jobs in the left-side navigation pane.
3. Click Ordinary Jobs on the Job Execution page.
4. On the Ordinary Jobs page, click History in the Actions column of an ordinary job. The Execution History page appears.

You can view the execution history of this job on the Execution History page. For more information, see [View the execution history](#).

3.1.13.2.3 Schemes

3.1.13.2.3.1 Create a scheme from a job

If an ordinary job or a cron job adapts to an O&M scenario of your product, you can save the job as a scheme to create product O&M tasks in the similar subsequent scenarios.

Prerequisites

You must have an ABM administrator account.

Context

Both cron jobs and ordinary jobs can be used to generate schemes. The procedures for these two types of jobs are the same. This topic uses the procedure for an ordinary job as an example.



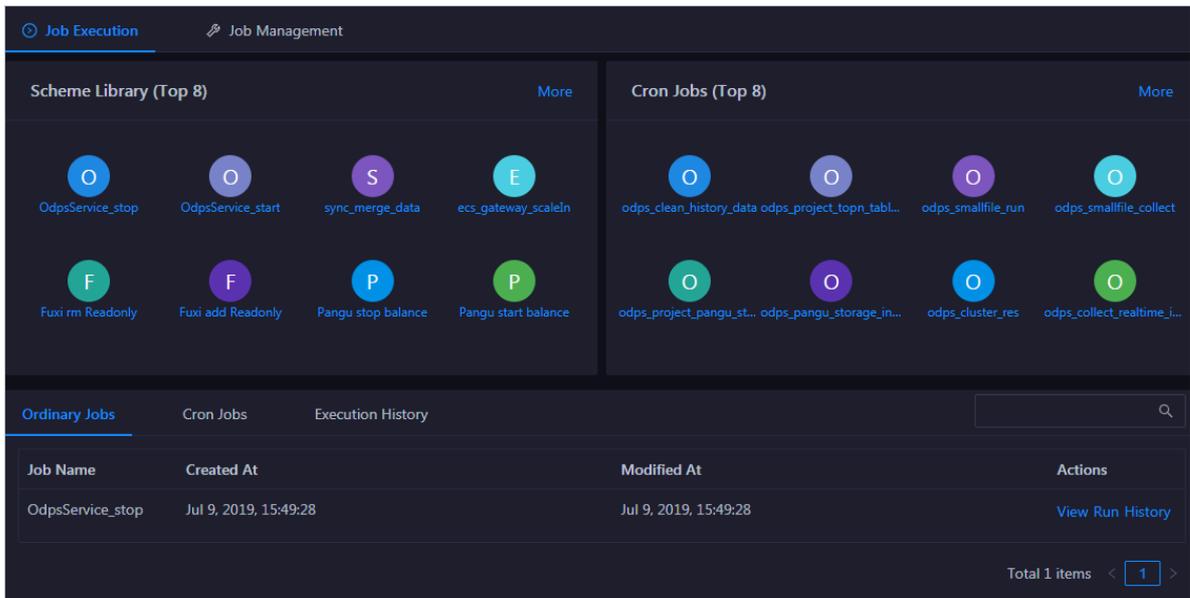
Notice:

When a cron job is saved as a scheme, no parameters are included.

Procedure

1. [Log on to the ABM console](#).
2. Click Management in the upper-right corner. On the page that appears, click Jobs in the left-side navigation pane.
3. Click Ordinary Jobs on the Job Execution page.

4. On the Ordinary Jobs page, click View in the Actions column of an ordinary job.



5. On the Job Details page, click Save as Scheme in the upper-right corner. The system prompts that you have saved the scheme.

Result

The new scheme has the same name as the job from which it was created and is listed on the Schemes page.

3.1.13.2.3.2 View schemes

The scheme is shown in the schemes list after it has been created. Apsara Bigdata Manager (ABM) allows you to view existing schemes in different ways, filter schemes, and search for specific schemes.

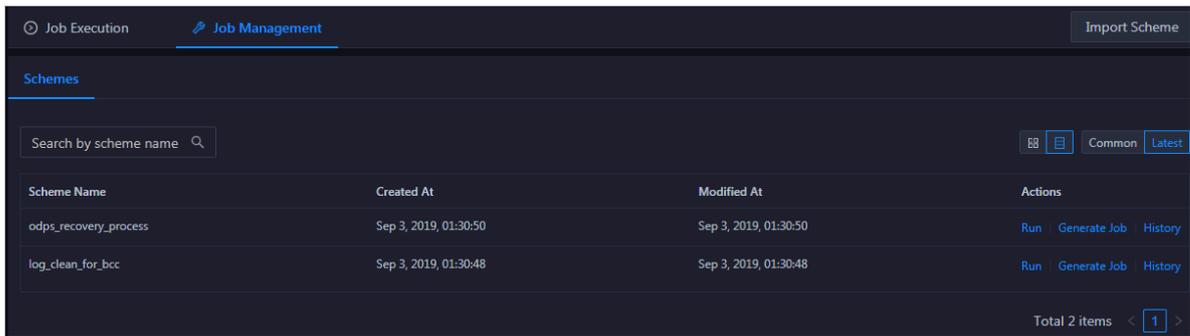
Prerequisites

You must have an ABM administrator account.

Procedure

1. [Log on to the ABM console.](#)
2. Click Management in the upper-right corner. On the page that appears, click Jobs in the left-side navigation pane.

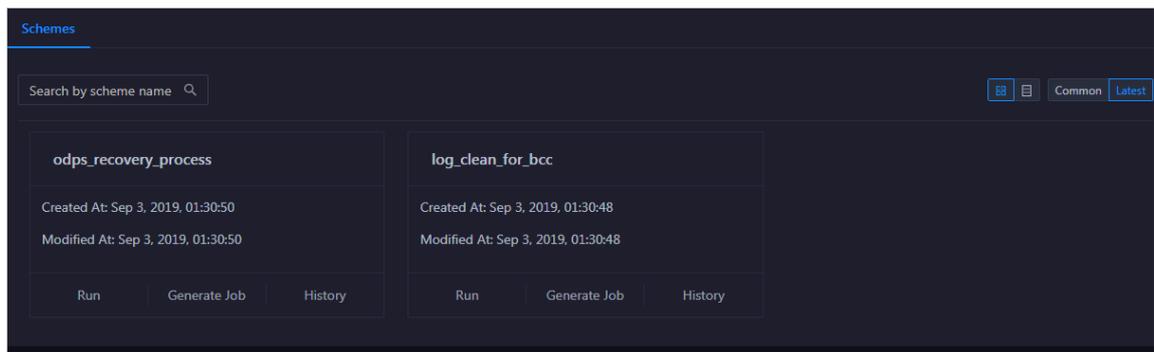
3. On the Jobs page, click Job Management.



4. If there are too many schemes, you can enter the scheme name in the search bar to search for the required scheme.

5. Change the method for viewing schemes.

- View schemes in list (default): Click  in the upper-right corner.
- View schemes in cards: Click  in the upper-right corner.



3.1.13.2.3 View the execution history of a scheme

Apsara Bigdata Manager (ABM) allows you to view the execution history of the required schemes to learn the execution status of them.

Prerequisites

You must have an ABM administrator account.

Procedure

1. [Log on to the ABM console.](#)
2. Click Management in the upper-right corner. On the page that appears, click Jobs in the left-side navigation pane.
3. On the Jobs page, click Job Management.

4. On the Schemes page, click History in the Actions column of a scheme that has directly run jobs.

You can view the execution history of this scheme on the Execution History page.

For more information, see [View the execution history](#).

3.1.13.2.4 View the execution history

Apsara Bigdata Manager (ABM) allows you to view the execution history of jobs and schemes so that you can learn about their execution details.

Prerequisites

You must have an ABM administrator account.

Context

After confirming the execution of a job, a record is automatically generated on the Execution History page.

The Execution History page provides the following features:

- Provides information such as the trigger mode, current status, start time, and end time of each job.
- Provides job execution details and parameter setting information, and allows you to download execution details.
- Allows you to perform the following operations depending on the job status. For example, you can run a job that is in the Pending status or retry the execution of a job that is in the Exception status.

Procedure

1. [Log on to the ABM console](#).
2. Click Management in the upper-right corner. On the page that appears, click Jobs in the left-side navigation pane.

3. Click the Execution History tab on the Job Execution page.

Job Name	Trigger Mode	Started At	Ended At	Status	Actions
odps_collect_realtime_instance_quota	Auto	Jul 7, 2019, 18:40:00	Jul 7, 2019, 18:40:07	Failure	View
odps_collect_project_meta	Auto	Jul 7, 2019, 18:40:00	Jul 7, 2019, 18:40:52	Success	View
odps_collect_cluster_quota_collect	Auto	Jul 7, 2019, 18:38:05	Jul 7, 2019, 18:38:16	Success	View
odps_collect_realtime_instance_quota	Auto	Jul 7, 2019, 18:38:00	Jul 7, 2019, 18:38:02	Failure	View
odps_collect_cluster_quota_collect	Auto	Jul 7, 2019, 18:36:05	Jul 7, 2019, 18:36:16	Success	View
odps_collect_realtime_instance_quota	Auto	Jul 7, 2019, 18:36:00	Jul 7, 2019, 18:36:01	Failure	View
odps_collect_cluster_quota_collect	Auto	Jul 7, 2019, 18:34:05	Jul 7, 2019, 18:34:16	Success	View
odps_collect_realtime_instance_quota	Auto	Jul 7, 2019, 18:34:00	Jul 7, 2019, 18:34:02	Failure	View

4. If there are too many execution records, filter them by a combination of one or more of the following filter conditions: job name, creator, execution status, and time range. Then, click to search for required records.

5. Click View in the Actions column of a record to view the execution details.

The following table lists the operations that you can perform on records in different statuses.

Execution status	Feature	Operation
All	View the parameter configuration	Click Parameter Configuration at the top, and select Context Parameters or Global Parameters to view the context parameters or global parameters of the task.
	Download execution details	Click Download Execution Details at the top to download the job execution details to the local device. Save it into a TXT file. The execution details record the JSON and raw data of job execution.

Execution status	Feature	Operation
	View execution details of steps	<ul style="list-style-type: none"> • On the Servers page of a step, click View Details in the Actions column of a certain server. The execution details of the step on the server, including the execution output, are displayed in the Execution Details section. • If the step includes a script, the Script Content and Execution Parameters pages will be displayed, where you can view the script content and the script execution parameters. • If the step includes a command, the Commands and Execution Parameters pages will be displayed, where you can view the command content and the command execution parameters.
	Refresh the page	If the task is in progress, you can click Refresh at the top to view the latest execution status.
Not Started	Start the execution	Click Start at the top to start the execution.
	Cancel the execution	Click Cancel at the top to cancel the execution.
Pending	Complete the manual operation and continue	At the manual step to be operated, follow the instructions and click OK to go to the next step.
	Roll back to the complete status of the previous step	At the manual step to be operated, click Rollback to roll back to the complete status of the previous step.
	Cancel the execution	Click Cancel to cancel the execution.
Exception	Retry the step with exception	At the step with exception, click Retry to execute the step again.
	Skip the step with exception	At the step with exception, click Skip to skip the step and continue.

Execution status	Feature	Operation
	Roll back to the complete status of the previous step	At the step with exception, click Rollback to roll back to the complete status of the previous step.
	Reset the step with exception to the Not Started status	<p>At the step with exception, click Reset to reset the step to the Not Started status.</p> <p>When the step with exception is reset to the Not Started status, the execution status becomes Paused. You can click Continue at the top to execute the step again.</p>
	View the execution details of exception steps	<ul style="list-style-type: none"> • On the Servers page of a step, click View Details in the Actions column of a certain server. The execution details of the step on the server, including the execution output and error message, are displayed in the Execution Details section. <p>After you have viewed the details of the server with exceptions during the execution, you can click Skip to skip this server.</p> <p>Alternatively, you can click Retry to execute the step again on the server.</p> <ul style="list-style-type: none"> • If the step includes a script, the Script Content and Execution Parameters pages will be displayed, where you can view the script content and the script execution parameters. • If the step includes a command, the Command Content and Execution Parameters pages will be displayed, where you can view the command content and the command execution parameters.
Failed	Retry the failed step	At the failed step, click Retry to execute the step again.
	Skip the failed step	At the failed step, click Skip to skip this step and execute the subsequent steps.

Execution status	Feature	Operation
	<p>Roll back to the complete status of the previous step</p>	<p>At the failed step, click Rollback to roll back to the complete status of the previous step.</p>
	<p>Reset the failed step to the Not Started status</p>	<p>At the failed step, click Reset to reset the step to the Not Started status.</p> <p>When the failed step is reset to the Not Started status, the execution status becomes Paused. Then, you can click Continue to execute the current step again.</p>
	<p>View execution details of failed steps</p>	<ul style="list-style-type: none"> • On the Servers page of a step, click View Details in the Actions column of a certain server. The execution details of the step on the server, including the execution output and error message, are displayed in the Execution Details section. After you have viewed the details of the server with exceptions during the execution, you can click Skip to skip this server. Alternatively, you can click Retry to execute the step again on the server. • If the step includes a script, the Script Content and Execution Parameters pages will be displayed, where you can view the script content and the script execution parameters. • If the step includes a command, the Command Content and Execution Parameters pages will be displayed, where you can view the command content and the command execution parameters.
	<p>Cancel the execution</p>	<p>Click Cancel at the top to cancel the execution.</p>

3.1.13.3 Patch management

Apsara Bigdata Manager (ABM) allows you to deploy and roll back upgrade patches for the products that it maintains. It also allows you to view detailed records of patch deployment and rollback by patch package or server.

Prerequisites

- **Your ABM account must have the required permissions to perform O&M operations on the corresponding product.**
- **You must have obtained the patch package in the `tar.gz` format for the product to be upgraded.**
- **The cluster of the product to be upgraded must be running properly.**

Entry

1. *Log on to the ABM console.*
2. **Click Management in the upper-right corner. On the page that appears, click Packages in the left-side navigation pane. The Packages page appears.**

Description of the Packages page:

- **Package Management:** allows you to manage the patch packages of the product. You can upload, deploy, or delete the packages.
- **Package Deployment:** displays the deployment history and details.

The Package Management page appears by default.

Upload a patch package

This section describes how to upload a patch package for ABM.

1. **Click Upload Package on the Package Management page.**
2. **In the dialog box that appears, select a patch package, and then click Upload. Wait until the uploading is complete.**

After the patch package is uploaded, the system prompts a success message. The patch package is then displayed in the list.

Deploy a patch package

After a patch package is uploaded, you can deploy it to the corresponding product cluster.

1. **In the patch package list, click Deploy in the Actions column of a patch package.**

2. In the dialog box that appears, set Cluster and Deployment Mode.

The valid values of Deployment Mode include:

- **All:** Deploy the patch package to all servers where it has not been deployed.
- **Phased Release:** Deploy the patch package on a random server.

3. Click OK.

The deployment status of the patch package is **Deploying**. Patch deployment takes some time. Wait until the patch package is deployed. Refresh the page after the deployment is complete. The deployment status is changed to **Deployed**.

Handle deployment failures

After you use ABM to deploy a patch for a product, the patch will be automatically bound to the service release (SR) version of the product. If the product is upgraded, the SR version is changed, and the deployment status of the patch package is changed to **Deployment Failed (Product Upgraded)**.

After the product is upgraded, ABM cannot determine whether the new version has fixed the problem to be resolved by the patch. Therefore, the patch automatically becomes invalid. If the product upgrade cannot fix the problem to be resolved by the patch, click **Force Deploy** to deploy the patch again. If the product upgrade has fixed the problem to be resolved by the patch, click **Ignore**.

View the deployment history and details

The **Deployment Records** page displays the deployment information about all patch packages. The **Deployment Details** page displays the deployment information about all servers.

1. Click the Package Deployment tab on the Packages page to view the deployment records.

The **Deployment Records** page displays the deployment records of all patch packages. You can view the name, version, product, cluster, service, service role, application type, deployment mode, and operation type of each patch package. You can also view the users who submitted the deployment requests, the total number of servers where each patch package needs to be deployed, the number of servers where each patch package is deployed, the number of servers

where each patch package fails to be deployed, the number of servers where the deployment has not finished, and the deployment time.

If there are too many deployment records, you can filter them by product name or package name.

2. Click the Deployment Details tab to view the deployment details.

The Deployment Details page displays the deployment information about all hosts, including the IP address, patch package name, version, product, cluster, service, service role, deployment progress, deployment status, associated build ID, deployment time, and log details.

If there are too many deployment details, you can filter them by product name, package name, or deployment status.

Roll back an upgrade patch

After an upgrade patch is deployed, you can roll back the cluster to the version before the deployment if the cluster runs abnormally or encounters other problems.

1. Click Roll Back in the Actions column of the patch package to be rolled back.



Note:

A patch package can be rolled back only when the deployment status is Deployed.

2. In the dialog box that appears, set Cluster to the cluster where the patch package is deployed, and then click OK.

Refresh the page during the rollback process. The deployment status is changed to Rolling Back. Rollback takes some time. Wait until the patch package is rolled back.

Refresh the page after the rollback is complete. The deployment status is changed to Rolled Back.

3.1.13.4 Hot upgrade

Apsara Bigdata Manager (ABM) allows you to upgrade monitoring configuration and metrics without interrupting the service. On the Hot Upgrades page, you can

view the hot upgrade history and upgrade logs. You can also delete the upgrade packages and upgrade history on this page.

Prerequisites

- **Your account must have the required permissions to perform O&M operations on ABM.**
- **You must have obtained the monitoring item upgrade package in the *tar.gz* format.**

Upgrade a monitoring item without interrupting the service

1. *Log on to the ABM console.*
2. **Click Management in the upper-right corner. On the page that appears, click Hot Upgrades in the left-side navigation pane.**
3. **Click Upload File, and then select and upload the obtained tar.gz file.**

The upload logs are displayed in the Upload Log section of the page during the upload process. After the upload is complete, the page displays the upgrade items for this upgrade package.

4. **Select the monitoring items to be upgraded, and then click OK.**
5. **In the dialog box that appears, click OK to start the upgrade.**

After the upgrade is complete, the system prompts that the upgrade is successful

.

View the hot upgrade history and logs

After the hot upgrade is complete, a hot upgrade record is generated on the File Management page, including the creation time and ID of the record, and the storage address of the upgrade package. When the hot upgrade fails, you can view the hot upgrade logs to locate the fault.

1. **Click the File Management tab on the Hot Upgrades page to view the hot upgrade history.**
2. **Click View Logs in the Actions column of an upgrade record to view the upgrade logs for each monitoring item in this hot upgrade process.**

Delete a hot upgrade record

ABM allows you to delete hot upgrade records, together with the corresponding hot upgrade packages and hot upgrade logs.

1. Click the File Management tab on the Hot Upgrades page to view the hot upgrade history.
2. Click Delete in the Actions column of an upgrade record. In the dialog box that appears, click OK.

3.1.13.5 Health management

Apsara Bigdata Manager (ABM) provides a wide range of built-in scheduling items and monitoring items for each product. These items check product faults and send alerts when necessary, enabling you to detect and fix product faults in time.

Prerequisites

- Your account must have the required permissions to perform O&M operations on the corresponding product.
- You must have obtained the alert sources and checkers of the monitoring items.

Background

Different products have different scheduling and monitoring items, but their configuration and operations are the same. This topic uses MaxCompute as an example.

Scheduling: You can run checkers on all hosts of a specified Apsara Infrastructure Management Framework role as scheduled to generate raw alert data. The raw alert data includes the checker, host, alert severity, and alert information. ABM stores the raw alert data in its database.

Monitoring: You can mount checkers to product pages in ABM. When mounting a checker to a product page, you can set a filter policy to display only required alerts.

Both the scheduling items and monitoring items are built-in and cannot be added. However, you can modify some parameters of the items, such as whether to enable an item, running parameters, and description. In addition, you can configure mount points of the monitoring items or delete monitoring items.

View details and mount points of scheduling items

The mount points of scheduling items are built-in and cannot be added, modified, or deleted. The mount points of the scheduling items correspond to the list of all servers corresponding to the Apsara Infrastructure Management Framework role that runs the scheduling script.

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click MaxCompute.
3. Click Management in the upper-right corner of the MaxCompute page, and then click Health Management in the left-side navigation pane of the Management page. The Scheduling page appears.

The Scheduling page displays all scheduling items of the current product.

4. On the Scheduling page, click View in the Actions column of a scheduling item to view the details.

The details of a scheduling item include the name, alias, description, alert cause, and alert solution.

5. Click + to expand a scheduling item, and then view the mount points of the scheduling item.

Modify a scheduling item

You can set the scheduling interval and running parameters of a scheduling item, and set whether to enable the scheduling item.

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click MaxCompute.
3. Click Management in the upper-right corner of the MaxCompute page, and then click Health Management in the left-side navigation pane of the Management page. The Scheduling page appears.
4. On the Scheduling page, click Edit in the Actions column of a scheduling item. In the dialog box that appears, set relevant parameters.

Type: The value System Default indicates that parameters such as Execution Interval and Parameters use the default settings. The value Custom indicates that the parameters can be customized.



Note:

Set the Execution Interval parameter based on the crontab command.

5. Click OK. The system prompts that the configuration has been modified.

View faulty servers

You can view all the faulty servers in the current cluster.

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click MaxCompute.
3. Click Management in the upper-right corner of the MaxCompute page, and then click Health Management in the left-side navigation pane of the Management page. The Scheduling page appears.
4. Click Faulty Servers in the upper-right corner to view the faulty servers in the cluster.

The faulty server list displays all faulty servers in the current cluster and the Apsara Infrastructure Management Framework role of each server.

Modify a monitoring item

You can modify the name and description of a monitoring item and determine whether to enable it. The alert sources and checkers of monitoring items are built-in. Do not modify them.

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click MaxCompute.
3. Click Management in the upper-right corner of the MaxCompute page, and then click Health Management in the left-side navigation pane of the Management page.
4. On the Health Management page, click the Monitoring tab. The Monitoring page appears.

The Monitoring page displays all monitoring items of the current product.

5. On the Monitoring page, click Modify in the Actions column of a monitoring item to modify its configuration.
6. Click OK. The system prompts that the configuration has been modified.

Add a mount point for a monitoring item

After a mount point is added for a monitoring item, the monitoring item mounts the raw alert data to the O&M page of each product in the ABM console.

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click MaxCompute.

3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the Management page.
4. On the Health Management page, click the **Monitoring** tab. The Monitoring page appears.
5. On the Monitoring page, click **+** to expand a monitoring item, and then view the mount points of the monitoring item.
6. Click **Add Mount Point** under the mount point list. In the dialog box that appears, set relevant parameters.

The following table describes some key parameters.

Parameter	Description
Mount Point	The mount point to which the required inspection result of this monitoring item is to be mounted. For example, the value odps/host indicates that the result is mounted to the host O&M page of MaxCompute.
Filter Policy	<p>Valid values:</p> <ul style="list-style-type: none"> · None: Display all alerts generated by the monitoring item. · Custom: Display the alerts generated by the monitoring item in accordance with the filter configured for the product tree node. · Node Name: Display the alerts whose node name is the same as the current node.
Enabled	Specifies whether the mount point takes effect.

7. Click **OK**. The system prompts that the configuration has been modified.

Delete a mount point for a monitoring item

1. [Log on to the ABM console](#).
2. Click  in the upper-left corner, and then click MaxCompute.
3. Click **Management** in the upper-right corner of the MaxCompute page, and then click **Health Management** in the left-side navigation pane of the Management page.
4. On the Health Management page, click the **Monitoring** tab. The Monitoring page appears.

5. On the Monitoring page, click + to expand a monitoring item, and then view the mount points of the monitoring item.
6. Click Delete in the Mount Point column of the mount point to be deleted. In the dialog box that appears, click OK. The system prompts that the deletion is successful.

Delete a monitoring item

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click MaxCompute.
3. Click Management in the upper-right corner of the MaxCompute page, and then click Health Management in the left-side navigation pane of the Management page.
4. On the Health Management page, click the Monitoring tab. The Monitoring page appears.
5. Click Delete in the Actions column of the monitoring item to be deleted. In the dialog box that appears, click OK. The system prompts that the deletion is successful.

3.1.13.6 Operation auditing

This feature allows you to view the O&M operations of the current product of Apsara Bigdata Manager (ABM). The details of each operation are provided for retrieval and fault locating.

Prerequisites

Your account must have the required permissions to perform O&M operations on the corresponding product.

Background

You can view operation logs by product. For example, to view the operation logs of MaxCompute, you must go to the MaxCompute page first. The following describes how to view the operation logs of MaxCompute.



Note:

This page displays only the O&M operations of a product. Note that the O&M operations of job services are not included.

Procedure

1. *Log on to the ABM console.*
2. Click  in the upper-left corner, and then click MaxCompute.
3. Click Management in the upper-right corner of the MaxCompute page, and then click Operation Audit in the left-side navigation pane of the Management page.

The Operation Audit page displays the O&M operations of the current product. In this example, the information about MaxCompute O&M operations is displayed, including the operation name, operation ID, status, submission time, start time, end time, operator, and implementation method.

4. Click Details for an operation to view the O&M operation details.

You can also view the causes of failed steps in detail.

5. If an O&M operation fails, view the cause of the failure.
6. When the task status is Failure, Not Started, Pending, or Exception, perform the operations listed in the following table based on your situation.

Status	Executable operation
Not Started	<ul style="list-style-type: none"> • Click Start to start the task. • Click Parameter Configuration to view the parameter configuration of the task. • Click Cancel to cancel the task.
Pending	<ul style="list-style-type: none"> • Follow the instructions and click OK to go to the next step. • Click Rollback to roll back to the complete status of the previous step. • Click Parameter Configuration to view the parameter configuration of the task. • Click Cancel to cancel the task.
Exception	<ul style="list-style-type: none"> • Click Retry to run the step again. • Click Skip to skip the step and continue. • Click Rollback to roll back to the complete status of the previous step. • Click Parameter Configuration to view the parameter configuration of the task. • Click Cancel to cancel the task.

Status	Executable operation
Failed	<ul style="list-style-type: none"> • Click Retry to run the step again. • Click Skip to skip the step and continue. • Click Rollback to roll back to the complete status of the previous step. • Click Parameter Configuration to view the parameter configuration of the task. • Click Cancel to cancel the task.

7. To download the O&M operation execution logs, click Download Execution Details at the top to save the logs to your local device.

3.1.14 Go to other platforms

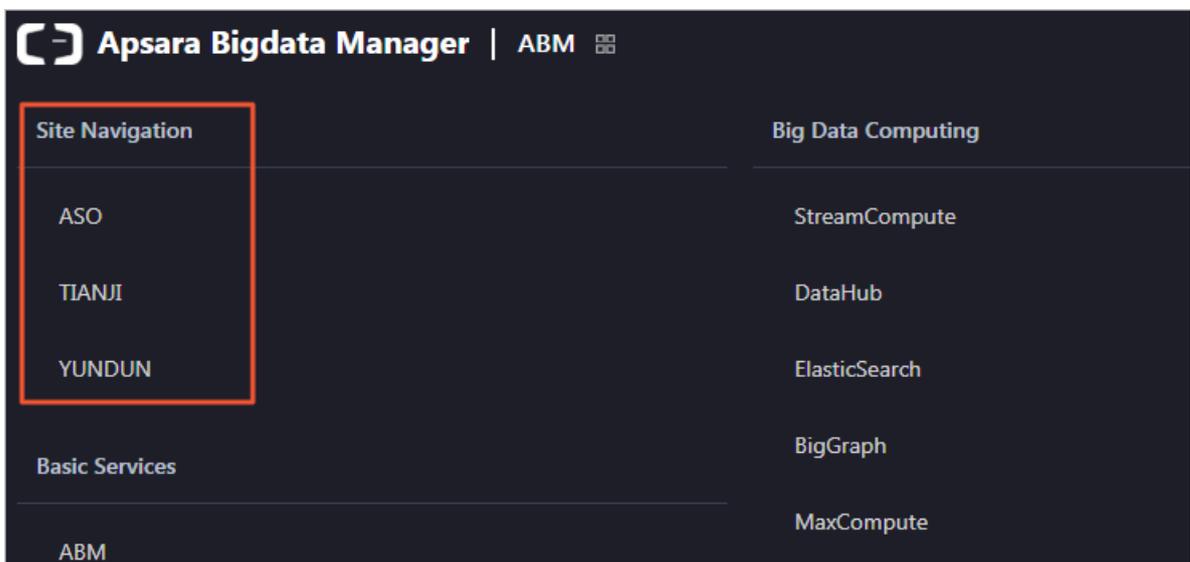
Apsara Bigdata Manager (ABM) provides the links to Apsara Stack Operation, Apsara Infrastructure Management Framework, and Alibaba Cloud Security to facilitate the O&M of big data products.

Prerequisites

You have obtained an ABM account that works properly and the corresponding password.

Procedure

1. *Log on to the ABM console.*
2. On the homepage of ABM, click  in the upper-left corner, and then click ASO, TIANJI , or YUNDUN in the Site Navigation section. The corresponding platform appears.



Result

After clicking ASO or TIANJI, you can log on to the corresponding platform without entering the username or password.

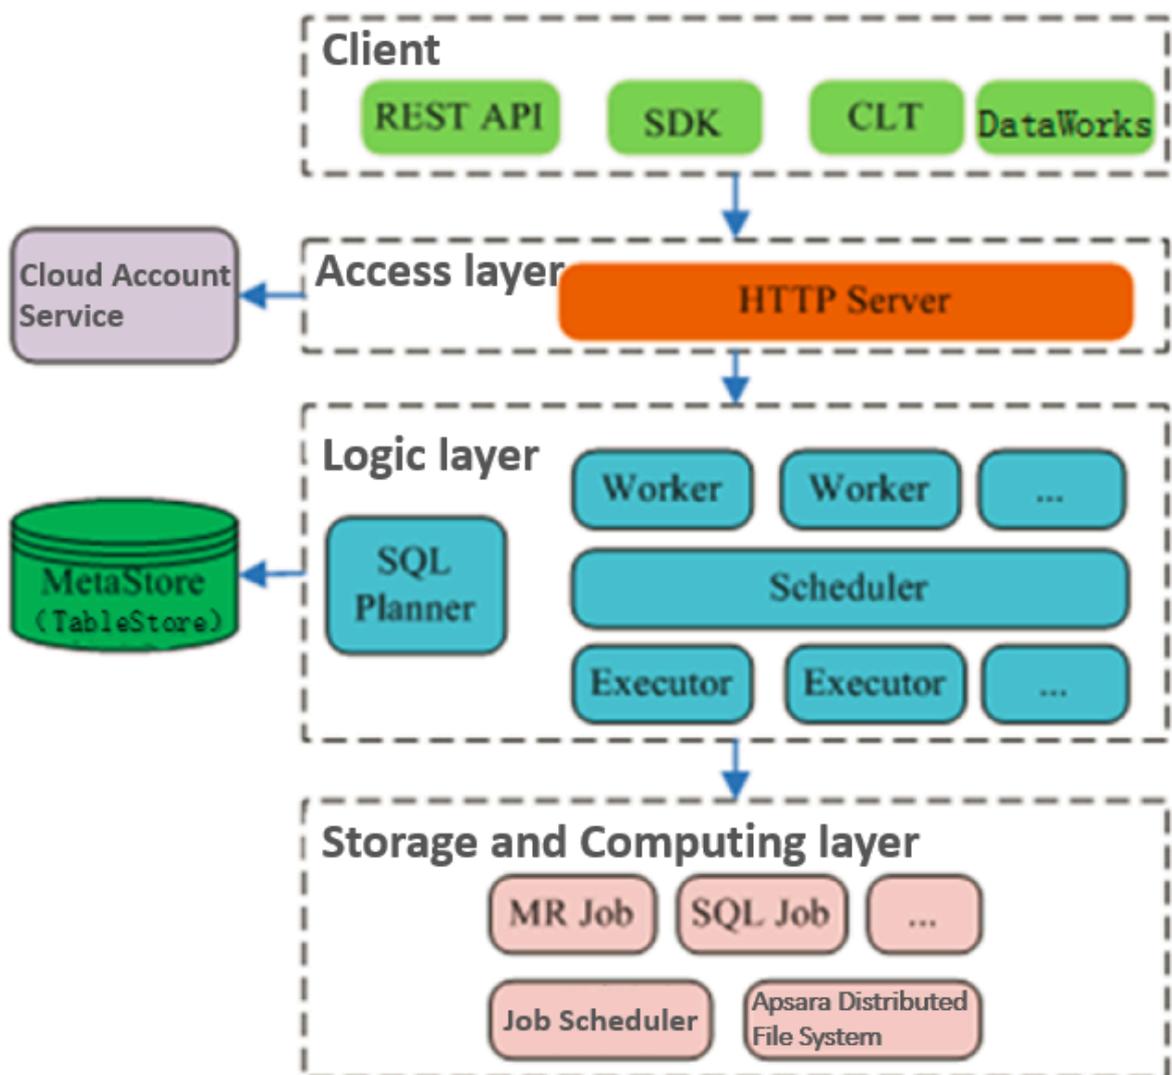
After clicking YUNDUN, however, you need to enter your username and password to log on to the platform.

3.2 MaxCompute

3.2.1 Concepts and architecture

Figure 3-4: MaxCompute architecture shows the MaxCompute architecture.

Figure 3-4: MaxCompute architecture



The MaxCompute service is divided into four parts: client, access layer, logic layer, and storage and computing layer. Each layer can be scaled out.

The following methods can be used to implement a MaxCompute client:

- **API:** RESTful APIs are used to provide offline data processing services.
- **SDK:** RESTful APIs are encapsulated in 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 data and use DDL and DML.
- **DataWorks:** provides upper-layer visual ETL and BI tools. DataWorks allows 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.

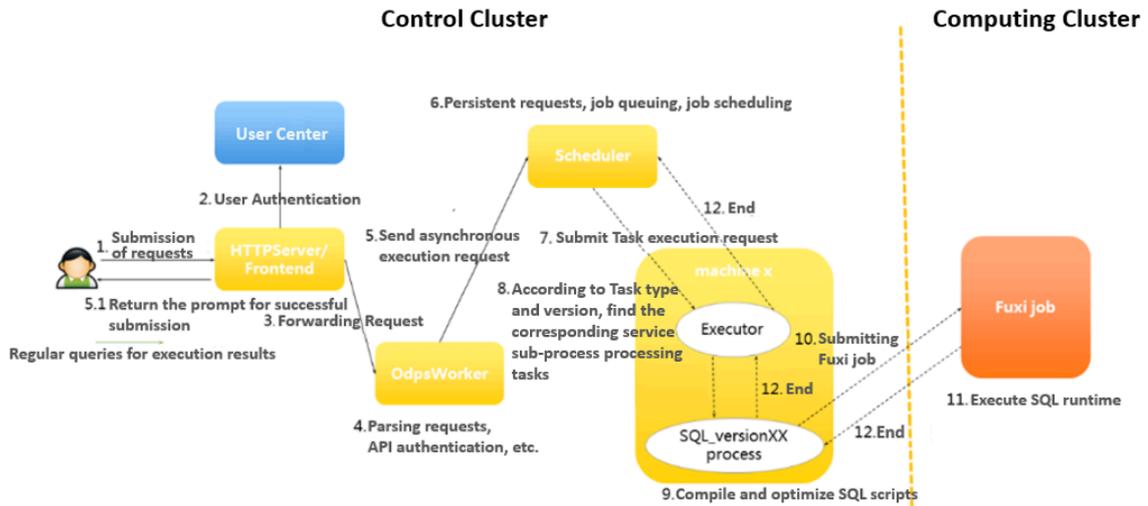
The logic layer is at the core of MaxCompute. It supports project and object management, command parsing and execution logics, and data object access control and authorization. The logic layer is divided into control and computing clusters. The control cluster is designed to manage projects and objects, parse and start queries and commands, and control and authorize access to data objects. The computing cluster executes tasks. Both control and computing clusters can be scaled in as needed. The control cluster includes three roles: Worker, Scheduler, and Executor. These roles are described as follows:

- **Worker** processes all RESTful requests, such as managing projects, resources and jobs. Worker forwards jobs that need to launch Fuxi tasks such as SQL, MapReduce, and Graph jobs to Scheduler for further processing.
- **Scheduler** schedules instances: splits an instance into multiple tasks, sorts tasks that are pending for submission, and queries resource usage from FuxiMaster in the computing cluster for throttling. If there are no idle slots in Job Scheduler, Scheduler stops processing task requests from Executor.
- **Executor** is responsible for launching SQL and MR tasks: submits Fuxi tasks to Fuxi Master in the computing cluster and monitors the operation of these tasks.

When you submit a job request, the Web server at the access layer queries IP addresses of registered Workers and sends API requests to randomly selected Workers. The Workers then send these requests to Scheduler for scheduling and throttling. Executor actively polls the Scheduler queue. If necessary resources

are available, it starts executing tasks and returns the task execution status to Scheduler. The following figure shows the MaxCompute job execution process.

Figure 3-5: MaxCompute job execution process



The MaxCompute job execution process involves the following concepts:

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 application is rarely implemented 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 close to 20 types of tasks, such as SQL, MapReduce, Admin, Lot, and Xlib. The execution logic of each task type varies greatly. Different tasks in an instance are identified by task_name. MaxCompute tasks run in the control cluster. Simple tasks such as metadata modification can run in the control cluster across their entire life cycles. To run computing tasks, submit Fuxi jobs to the computing cluster.

3. **Fuxi job:** a computing model provided by the Job Scheduler module corresponding 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 jobs and online jobs (which evolve from the service modes). An online job is also called a quasi-real-time task. An online job is a resident process that can be executed any time there are tasks, reducing the time required to start and stop a job.
 - You can submit a MaxCompute task to multiple computing 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 logics.
5. **Fuxi instance:** the instance of a Fuxi task. Fuxi instances are the smallest units 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 re-used 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 fault locating. 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 life cycles.

The storage and computing layer of MaxCompute is a core component of Alibaba Cloud proprietary cloud computing platform. As the kernel of the Apsara system

, this layer runs in the computing cluster independent of the control cluster. The architecture diagram illustrates only the major modules.

3.2.2 O&M commands and tools

3.2.2.1 Before you start

Before using MaxCompute O&M commands and tools, you must be aware of the following information:

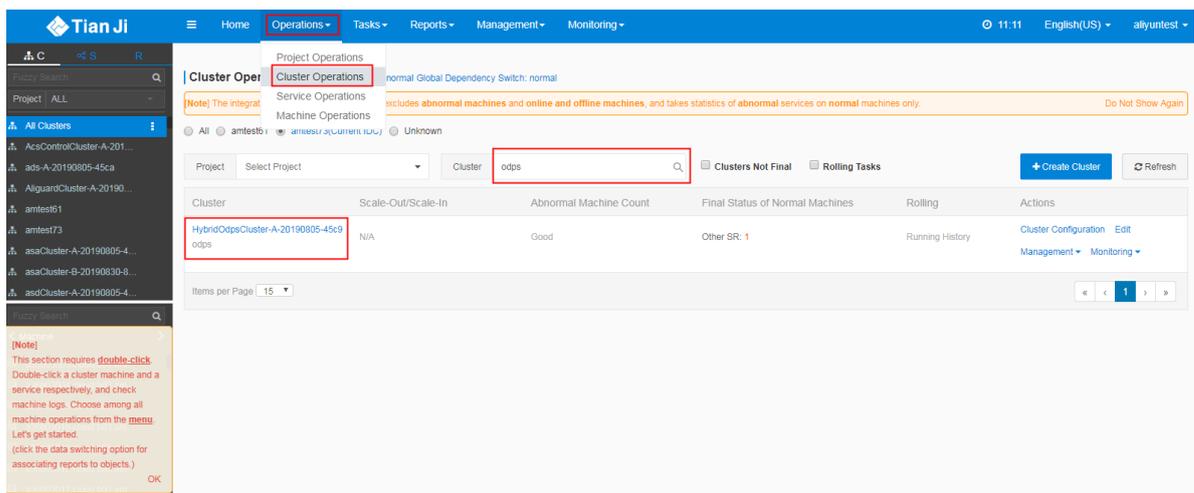
During the MaxCompute O&M process, the default account is admin. You must run all commands as an admin user. You must use your admin account and sudo to run commands that require sudo privileges.

3.2.2.2 odpscmd commands

You can perform operations and maintenance through command lines. You need to log on to the command line tool before using the commands. The specific procedure is as follows:

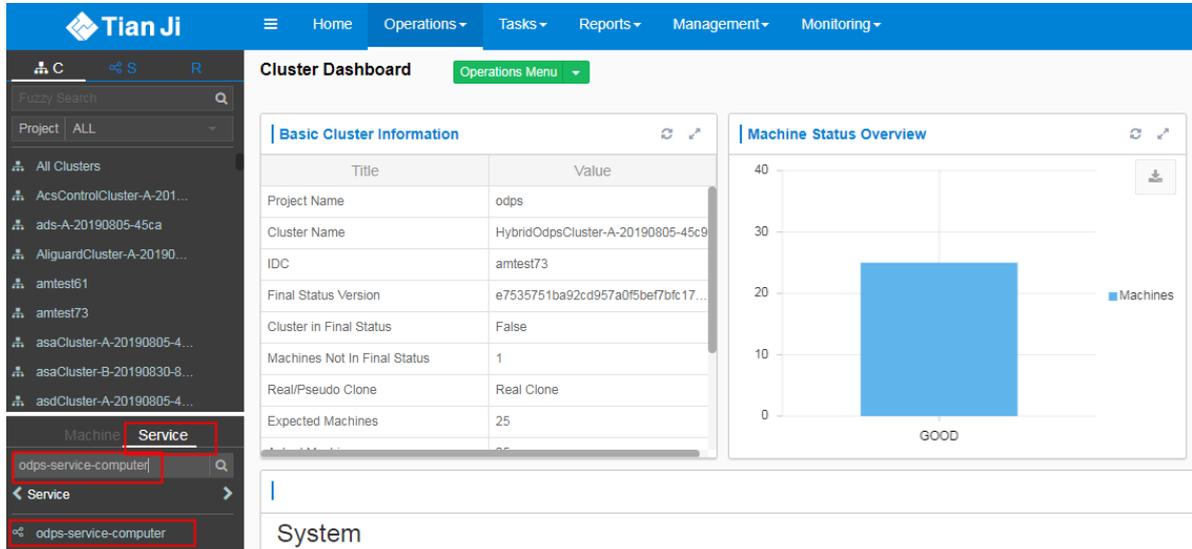
1. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. In the Cluster search box, enter **odps** to search for the expected cluster.

Figure 3-6: Search for a cluster



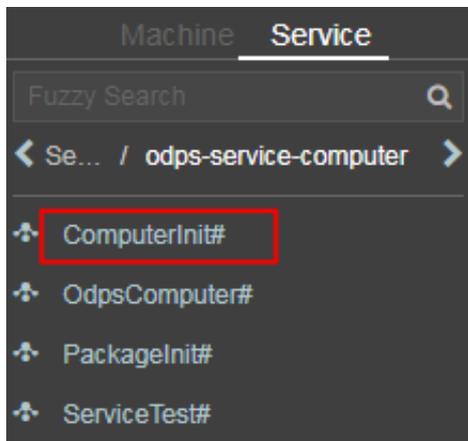
2. Click the cluster in the search result. In the left-side navigation pane, click the Service tab, and locate and double-click the odps-service-computer service.

Figure 3-7: odps-service-computer



3. After you access the odps-service-computer service, double-click ComputerInit#.

Figure 3-8: ComputerInit#



4. On the tab page that appears, hover over the vertical dots and choose Terminal from the shortcut menu. In the TerminalService window that appears, you can perform subsequent command line operations.

Figure 3-9: Terminal menu

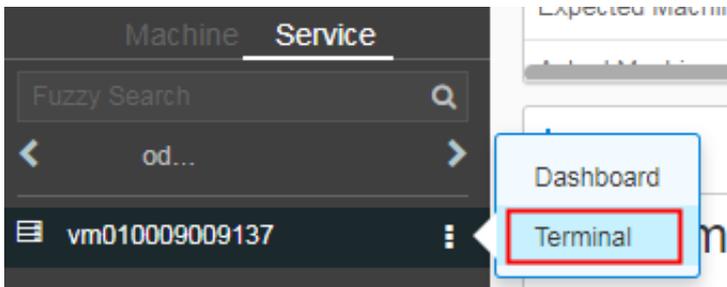
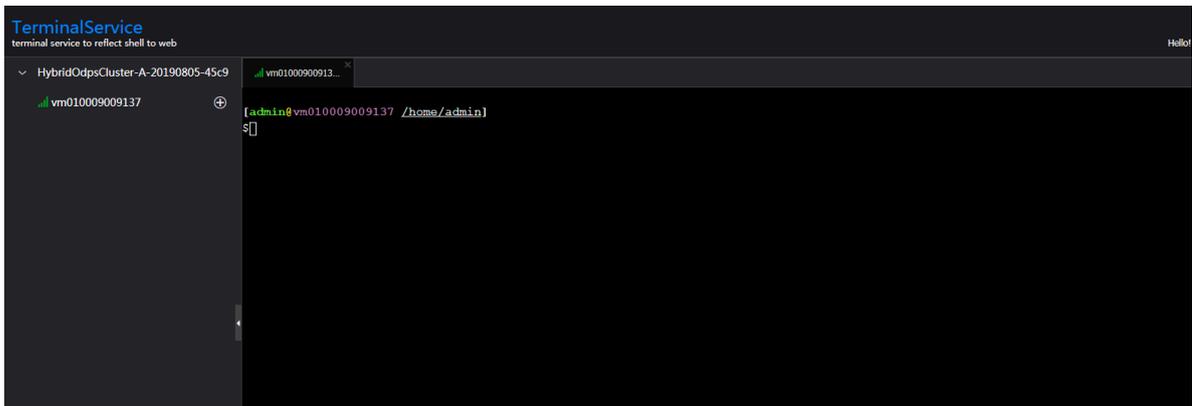


Figure 3-10: TerminalService command line window



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 ID, Key, `end_point`, `log_view`, and `tunnel_point` are configured by default. You can directly use the `./clt/bin/odpscnd` command to view information such as the version number in interactive mode. For example, you can run the `HTTP GET /projects/admin_task_project/system;` command to check the version information of MaxCompute.

Description of client command parameters

The following figure shows the client command parameters.

Figure 3-11: Client command parameters

```

$ /apsara/odps_tools/ctl/bin/odpscmd -h
Usage: odpscmd [OPTION]...
where options include:
  --help                (-h) for help
  --config=<config_file>  specify another config file
  --project=<prj_name>    use project
  --endpoint=<http://host:port>  set endpoint
  -u <user_name> -p <password>  user name and password
  --instance-priority=<priority>  priority scope[0-9]
  -M                    read machine readable data
  -k <n>                will skip beginning queries and start from specified position
  -r <n>                set retry times
  -f <"file_path;">    execute command in file
  -e <"command;[command;]...">  execute command, include sql command
  -C                    will display job counters
  -V                    will not submit jobs to fuxi master
    
```

- **-e:** The MaxCompute client does not execute SQL statements in interactive mode.
- **--project, -u, and -p:** The client directly uses the values specified for the project, user, and pass parameters. When you specify values for parameters such as user and pass, the client uses the specified values, instead of the values configured in the conf file. For other parameters, the client uses the values configured in the conf file.
- **-k and -f:** The client directly executes local SQL files.
- **--instance-priority:** This parameter is used to assign a priority to the current task. Valid values: 0 to 9. A lower value indicates a higher priority.
- **-r:** This parameter indicates the number of retries. It is commonly used in scripting jobs.

Commonly used SQL commands for O&M

The following table lists the commonly used commands.

Table 3-8: Commonly used commands

Command	Description
whoami;	Allows you to view your Alibaba Cloud account and endpoint information.
show p;	Allows you to view information about all instances that have been run.

Command	Description
wait <instanceid>;	Allows you to re-generate the log information and Fuxi job information of a task. To do so, you must have the owner permission, and the log information and Fuxi job information must be in the same project.
kill <instanceid>;	Allows you to terminate specified instances.
tunnel upload/download;	Allows you to test whether Tunnel functions properly.
desc project <projectname> -extended;	Allows you to view the project usage. <ul style="list-style-type: none"> • 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;	Allows you to export DLL statements of all tables in a project.
create table tablename (...);	Allows you to create a table.
select count(*) from tablename;	Allows you to search for a table.
Explain	Allows you to create plans without submitting Fuxi jobs and view resources required for the jobs.
list	Allows you to list tables, resources, and roles.
show	Allows you to view table and partition information.
purge	Allows you to remove all data from the MaxCompute recycle bin directly to the Apsara Distributed File System recycle bin. <ul style="list-style-type: none"> • purge table <tablename>: allows you to purge a single table. • purge all: allows you to purge all tables in the current project.

3.2.2.3 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.

Table 3-9: 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]>
       upload data from local file
  -acp,-auto-create-partition <ARG>  auto create target partition if
not                                     exists, default false
  -bs,-block-size <ARG>              block size in MiB, default 100
```

```

-c,-charset <ARG>          specify file charset, default
ignore.                    ignore.
                             set ignore to download raw data
                             compress, default true
-cp,-compress <ARG>        specify discard bad records
-dbr,-discard-bad-records <ARG> action(true|false), default false
                             specify date format pattern,
-dfp,-date-format-pattern <ARG> default
                             yyyy-MM-dd HH:mm:ss
                             specify field delimiter, support
-fd,-field-delimiter <ARG> unicode, eg \u0001. default ",",
                             if local file should have table
-h,-header <ARG>           header, default false
                             max bad records, default 1000
-mbr,-max-bad-records <ARG> specify null indicator string,
-ni,-null-indicator <ARG> default ""(empty string)
                             specify record delimiter, support
-rd,-record-delimiter <ARG> unicode, eg \u0001. default "\r\n"
"
-s,-scan <ARG>             specify scan file
                             action(true|false|only), default
true                        true
-sd,-session-dir <ARG>     set session dir, default
                             D:\software\odpscmt_public\
                             plugins\ds
                             hip
-ss,-strict-schema <ARG>   specify strict schema mode. If
false,                      extra data will be abandoned and
                             insufficient field will be filled
                             with null. Default true
                             tunnel endpoint
-te,-tunnel_endpoint <ARG> tunnel endpoint
                             number of threads, default 1
-threads <ARG>             time zone, default local timezone
-tz,-time-zone <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 download [options] <[project.]table[/partition]> <path>
       download data to local file
  -c,--charset <ARG>          specify file charset, default
ignore.
  -ci,--columns-index <ARG>   set ignore to download raw data
from                          specify the columns index(starts
each                           0) to download, use comma to split
                               index
  -cn,--columns-name <ARG>    specify the columns name to
download,                      use comma to split each name
  -cp,--compress <ARG>       compress, default true
  -dfp,--date-format-pattern <ARG> specify date format pattern,
default                        yyyy-MM-dd HH:mm:ss
  -e,--exponential <ARG>     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
                               unicode, eg \u0001. default ","
  -h,--header <ARG>          if local file should have table
header,                        default false
  -limit <ARG>                specify the number of records to
                               download
  -ni,--null-indicator <ARG> specify null indicator string,
default                        ""(empty string)
  -rd,--record-delimiter <ARG> specify record delimiter, support
                               unicode, eg \u0001. default "\r\n"
  -sd,--session-dir <ARG>    set session dir, default
                               D:\software\odpscmd_public\plugins\
dshi
p
```

```

-te,-tunnel_endpoint <ARG>      tunnel endpoint
  -threads <ARG>                 number of threads, default 1
-tz,-time-zone <ARG>           time zone, default local timezone:
                                Asia/Shanghai
usage: tunnel download [options] instance://<[project/]instance_id> <
path>
                                download instance result to local file
  -c,-charset <ARG>             specify file charset, default
ignore.
  -ci,-columns-index <ARG>      set ignore to download raw data
from                                specify the columns index(starts
each                                from
                                0) to download, use comma to split
  -cn,-columns-name <ARG>       index
download,                          specify the columns name to
  -cp,-compress <ARG>           use comma to split each name
  -dfp,-date-format-pattern <ARG> compress, default true
default                             specify date format pattern,
  -e,-exponential <ARG>        yyyy-MM-dd HH:mm:ss
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
unicode, eg \u0001. default ","
  -h,-header <ARG>             if local file should have table
header,                             default false
  -limit <ARG>                 specify the number of records to
download
  -ni,-null-indicator <ARG> specify null indicator string, default
""(empty string)
  -rd,-record-delimiter <ARG> specify record delimiter, support
unicode, eg \u0001. default "\r\n"
  -sd,-session-dir <ARG>      set session dir, default
D:\software\odpscmd_public\plugins\
dshi
  -te,-tunnel_endpoint <ARG>   p
  -threads <ARG>               tunnel endpoint
  -tz,-time-zone <ARG>         number of threads, default 1
                                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.

**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
```

3.2.2.4 LogView tool

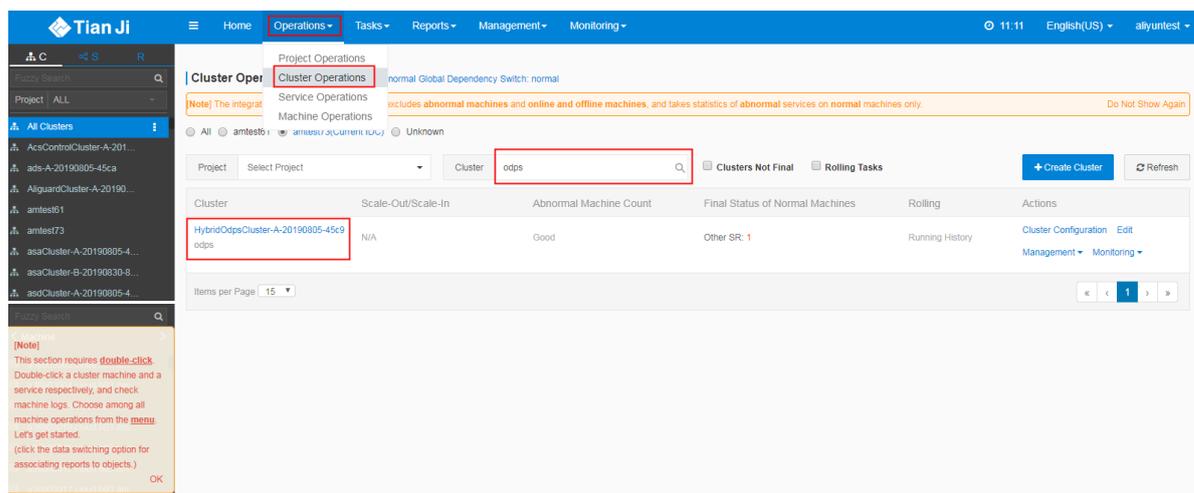
3.2.2.4.1 Before you start

You must confirm the LogView process status before using LogView. If the process status is off, you need to start the LogView process.

The procedure for querying the process status and starting the process is as follows:

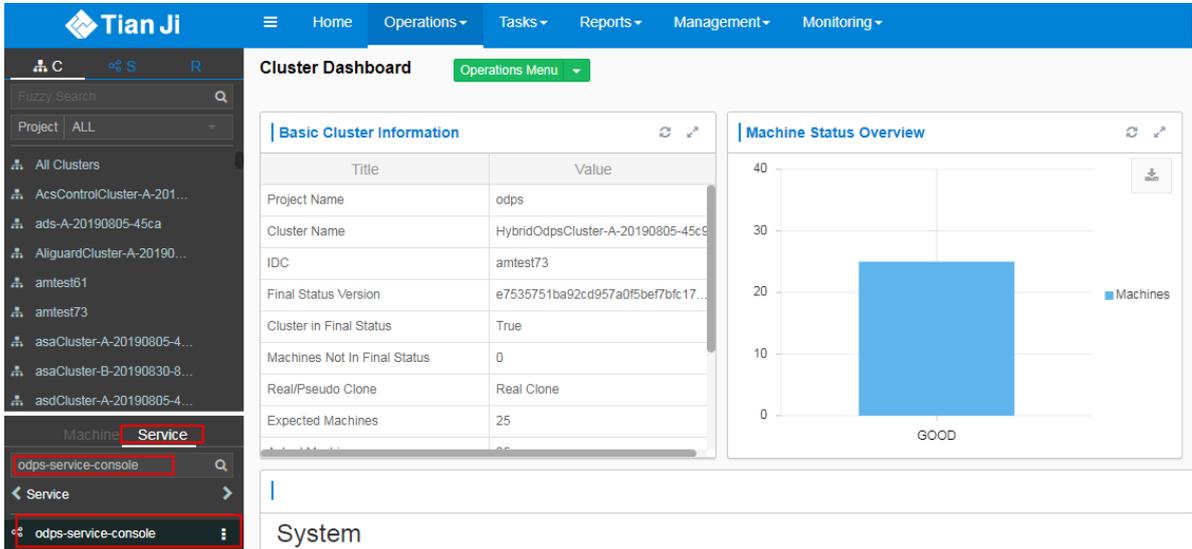
1. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. In the Cluster search box, enter **odps** to search for the expected cluster.

Figure 3-12: Search for a cluster



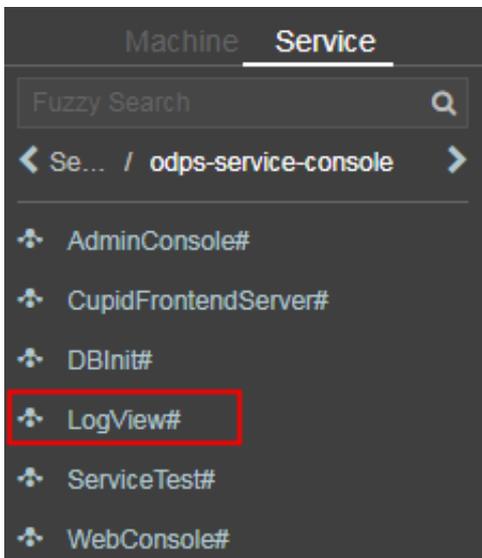
2. Click the cluster in the search result. In the left-side navigation pane, click the Service tab, and locate and double-click the odps-service-computer service.

Figure 3-13: odps-service-console service



3. After you access the odps-service-console service, double-click LogView#.

Figure 3-14: LogView#



4. On the tab page that appears, hover over the vertical dots and choose **Terminal** from the shortcut menu to open the TerminalService window.

Figure 3-15: Terminal menu

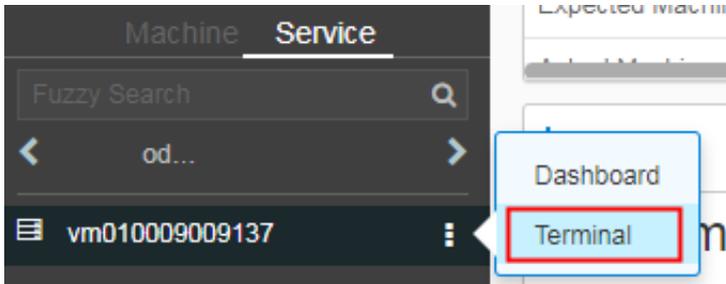
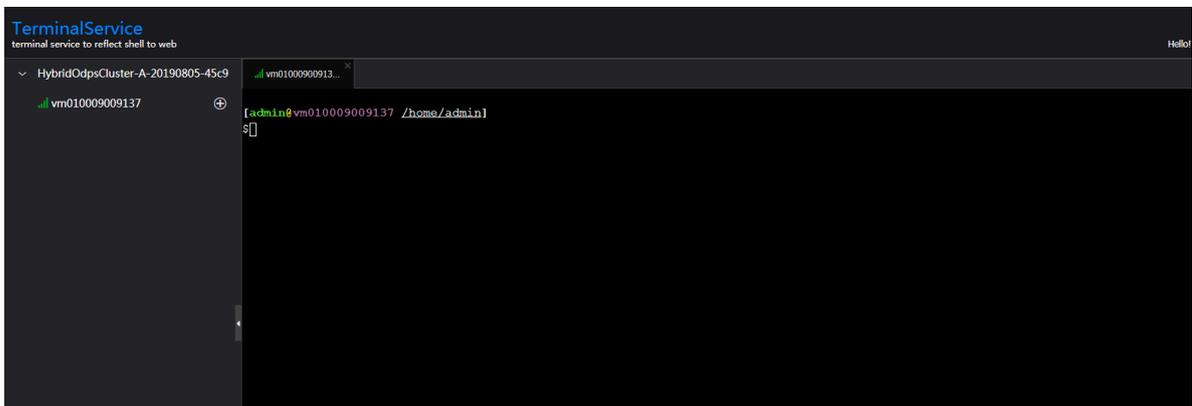


Figure 3-16: TerminalService command line window



5. Run the following command to locate the Docker container where LogView resides:

```
docker ps|grep logview
```

6. Run the following command to view the LogView process status:

```
ps -aux|grep logview
```

```
netstat -ntulp|grep 9000
```

7. 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.

3.2.2.4.2 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.

Figure 3-17: A long string starting with logview

```
ID = 20151214065043617g1jgn2i8
log view:
http://logview.odps.aliyun.com/logview/?h=http://service.odps.aliyun.com/api&p=yunxiang_01&i=20151214065043617g1jgn2i8&token=NTA2QDA2NDMseyJTdGF0ZW11bn0iOi0t7IkFjdG1vb1I6WyJvZHBzO1JlYWQiXSwiRWZmZWNOIjoiQWxsY3ciLCJSZXNvdXJjZSI6WyJhY3M6b2RwczoqOnByb2VmYyc2Ivbi16IiEif0==
```

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.

Figure 3-18: Composition of a LogView string

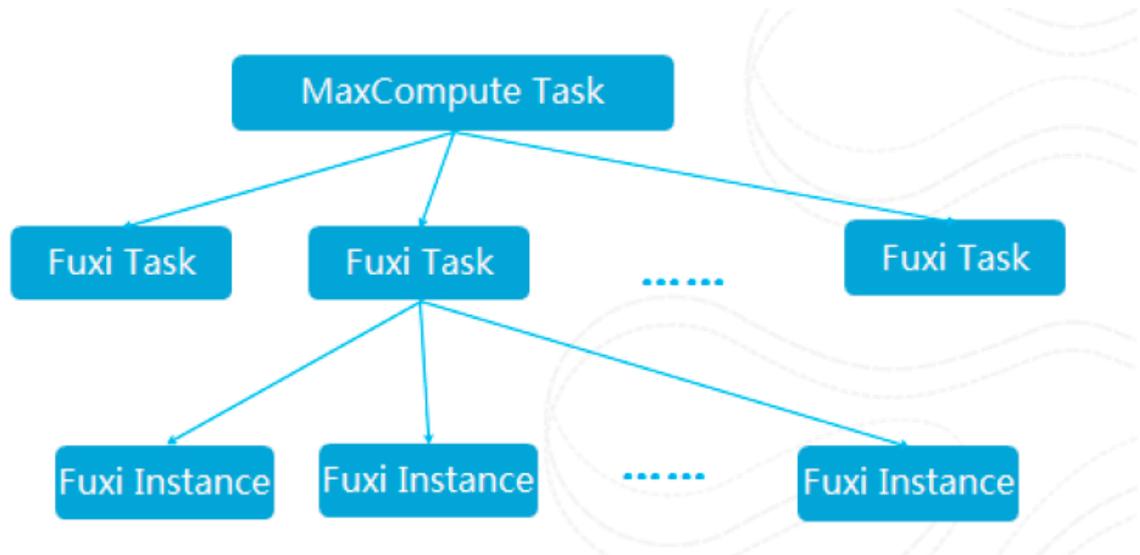
```
http://logview.odps.aliyun.com/logview/
?h=http://service.odps.aliyun.com/api
api&p=yunxiang_01
&i=20151214065043617g1jgn2i8
&token=WGhVU2haQXNha0t1V0FOWlRPLzZWk3hPMXFVPSxPRFE
```

3.2.2.4.3 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.

Figure 3-19: Relationships between MaxCompute tasks and Fuxi instances



The following figures show the relevant information in LogView.

MaxCompute Instance

Figure 3-20: MaxCompute Instance

URL	Project	InstanceID	Owner	StartTime	EndTime	Status	SourceXML
http://service.odps.aliyun.co...	yunxiang_01	20151214065043617g...	ALYINJtrain...	2015-12-14 14:5...	2015-12-14 14:5...	Terminated	SQL

```

Node XML: [console_select_query_task_1450075843613]
<SQL>
<Name>console_select_query_task_1450075843613</Name>
<Config>
<Property>
<Name>settings</Name>
<Value>{"odps.idata.useragent":"CLT(0.17.3 : 9a2149c); Windows 7(10.10.52.38/ali-87315n)","odps.sql.select.output.format":"HumanReadable"}
</Value>
</Property>
<Property>
<Name>guid</Name>
<Value>69f56821-a782-45b6-9668-34a7eb4ed5d6</Value>
</Property>
<Property>
<Name>uuid</Name>
<Value>46c46f5d-cb0b-4b74-9d2d-a32e64e63dd8</Value>
</Property>
</Config>
<Query>select count(*) from t.test.ni;</Query>
  
```

```

Source for: 20151214065043617g1jgn2i8
<?xml version="1.0" encoding="UTF-8"?>
<Job>
<Priority>9</Priority>
<Tasks>
<SQL>
<Name>console_select_query_task_1450075843613</Name>
<Config>
<Property>
<Name>settings</Name>
<Value>{"odps.idata.useragent":"CLT(0.17.3 : 9a2149c); Windows 7(10.10.52.38/87315n)","odps.sql.select.output.format":"HumanReadable"}</Value>
</Property>
<Property>
<Name>guid</Name>
<Value>69f56821-a782-45b6-9668-34a7eb4ed5d6</Value>
</Property>
<Property>
<Name>uuid</Name>
<Value>46c46f5d-cb0b-4b74-9d2d-a32e64e63dd8</Value>
</Property>
</Config>
<Query>select count(*) from t.test.ni;</Query>
</SQL>
  
```


Task Detail - Fuxi Job

Figure 3-22: Task Detail - Fuxi Job(1)

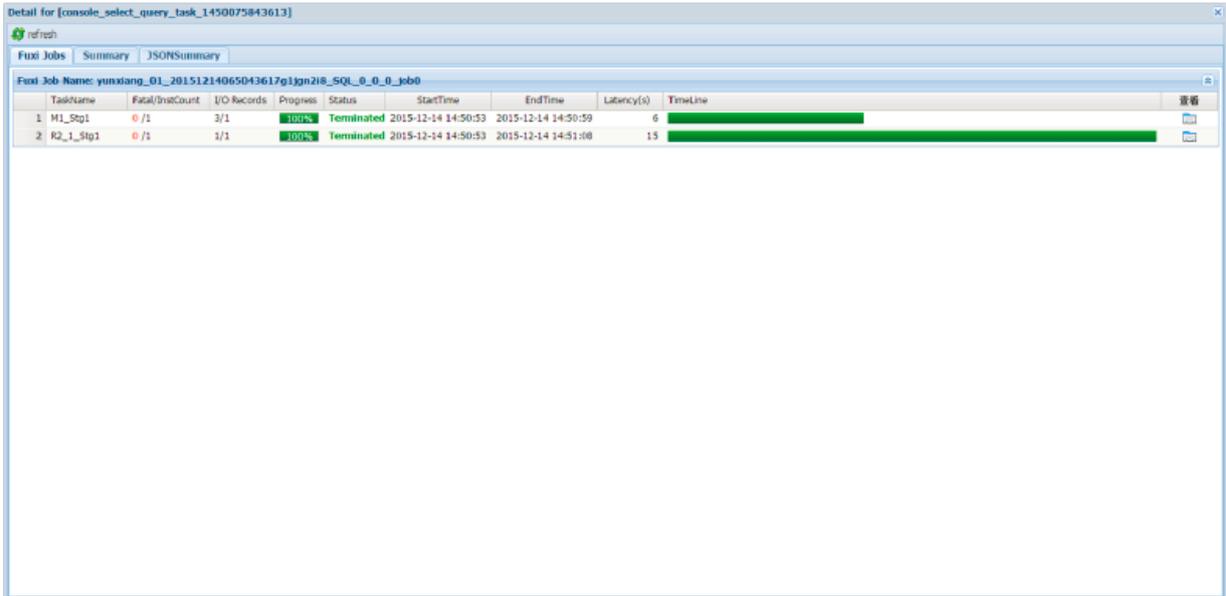
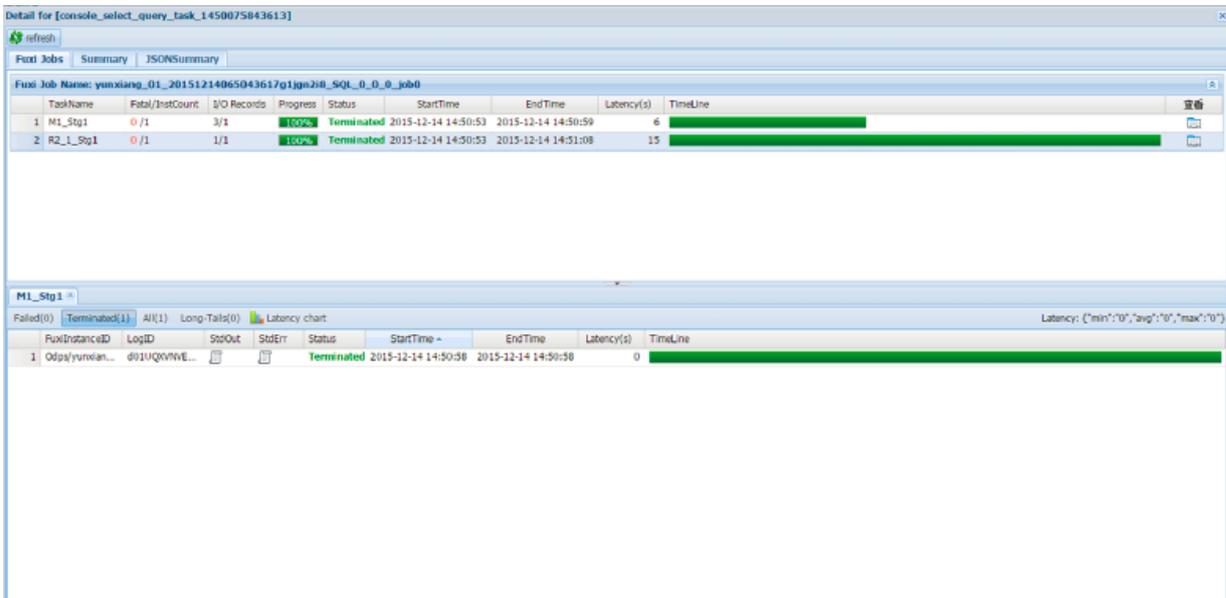


Figure 3-23: Task Detail - Fuxi Job(2)



Task Detail - Summary

Figure 3-24: Task Detail - Summary

Detail for [console_select_query_task_1450075843613]

refresh

Fuxi Jobs Summary JSONSummary

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)

```
Summary:
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

Figure 3-25: Task Detail - JSONSummary

```

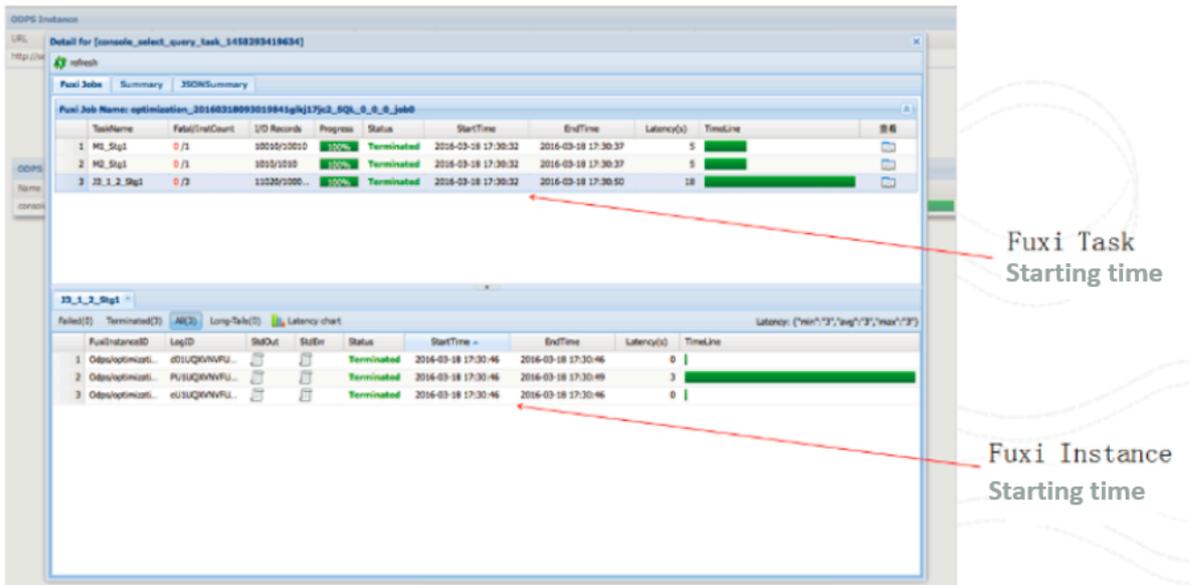
Detail for [console_select_query_task_1450075843613]
refresh
Fuxi Jobs Summary JSONSummary
{
  "inputs":
  {
    "yunxiang_01.t_test_ni":
    [
      3,
      824
    ]
  },
  "jobs":
  [
    {
      "inputs":
      {
        "M1_Stg1": 3
      },
      "tasks":
      {
        "M1_Stg1":
        {
          "writer_dumps":
          {
            "R2_1_Stg1":
            [
              0,
              0,
              0
            ]
          }
        },
        "planned_memory": 2048,
        "input_attrs":
        {
          "split": 256
        },
        "user_counters":
        {
        },
        "total_instance_run_time": 0,
        "output_record_counts":
        {
          "R2_1_Stg1":
          {
            "R2_1_Stg1": 1
          }
        }
      }
    }
  ]
}

```

3.2.2.4.4 Basic operations and examples

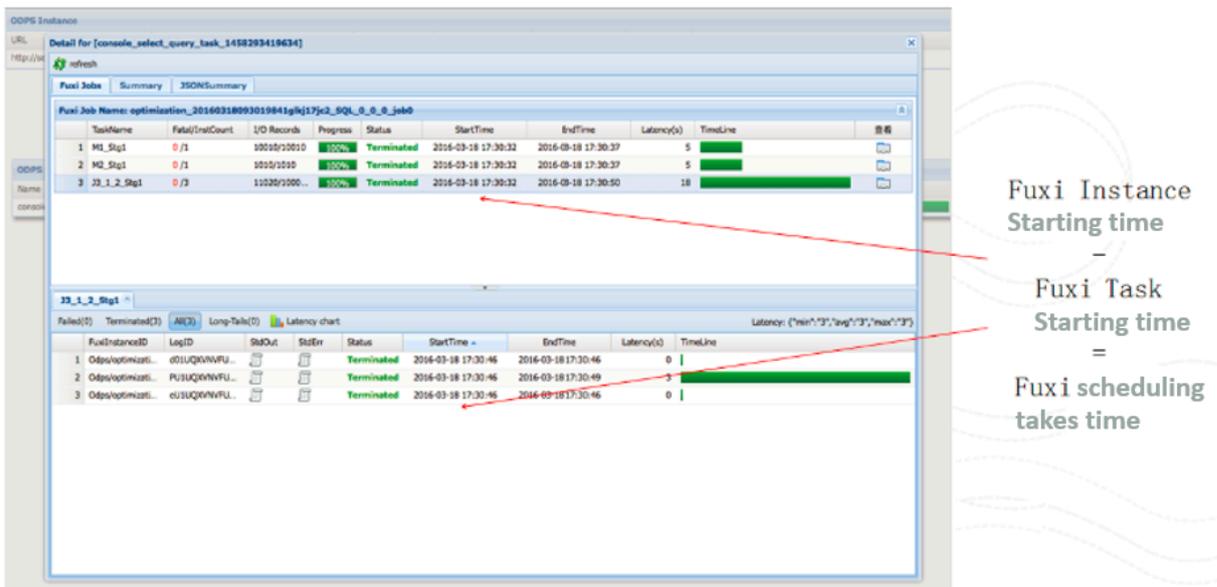
View each point in time in the life cycle of a job.

Figure 3-26: View each point in time in the life cycle of a job



View the time it takes for Job Scheduler to schedule an instance.

Figure 3-27: View the time it takes for Job Scheduler to schedule an instance



View the polling interval.

Figure 3-28: View the polling interval

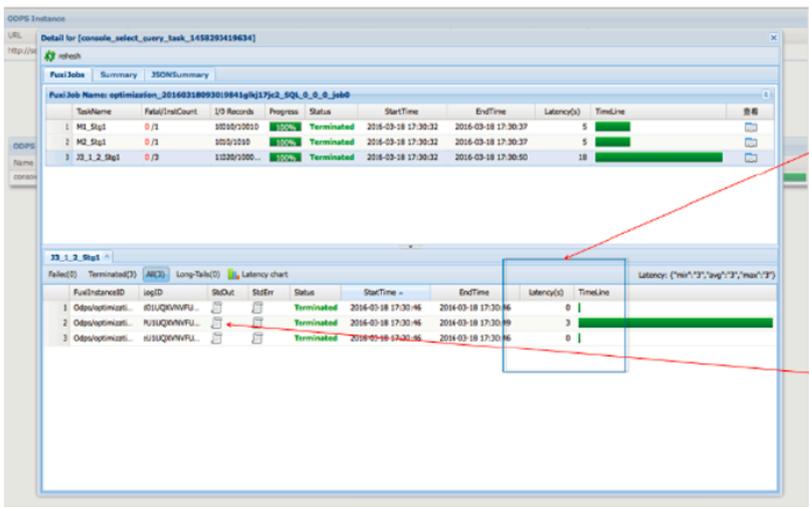
```
odps@ optimization>select * from skew a join small b on a.key=b.key;

ID = 20160318092653630gstax6jc2
Log view:
http://logview.odps.aliyun.com/logview/?h=http://service.odps.aliyun.com/api&p=optimization&i=20160318092653630gstax6jc2&token=d05vbmZOWUpSRkhCVllzUhdGM3I1SEFoeEFVPSxPRFBTX09CTzoxMDExODIyNTI0ODIzNDU5LDE0NTg4OTgwMTMseyJTdGF0ZW1lbnQiOiI7IkFjdGlvbiI6WyJvZHBzO1JlYWQiXSwiRWZmZWNOIjoiQWxsY3c1L0JlZSZNvZjZjZSI6WyJhY3M6b2RwczoqOnByb2p1Y3RzL29wdGltZXphdGlvbi9pbmN0YW5jZXMvMjAxNjAzMTgwOTIzNTMzMzBnc3RhedZqYzIiXX1dLCJWZXJzaW9uIjojMSJ9
2016-03-18 17:27:05 M1_Stg1_job0:0/0/1[0%] M2_Stg1_job0:0/0/1[0%] J3_1_2_Stg1_job0:0/0/3[0%]
2016-03-18 17:27:10 M1_Stg1_job0:0/1/1[100%] M2_Stg1_job0:0/1/1[100%] J3_1_2_Stg1_job0:0/0/3[0%]
2016-03-18 17:27:16 M1_Stg1_job0:0/1/1[100%] M2_Stg1_job0:0/1/1[100%] J3_1_2_Stg1_job0:0/0/3[0%]
Summary:
resource cost: cpu 0.02 Core * Min, memory 0.03 GB * Min
```

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

Figure 3-29: Check for data skews

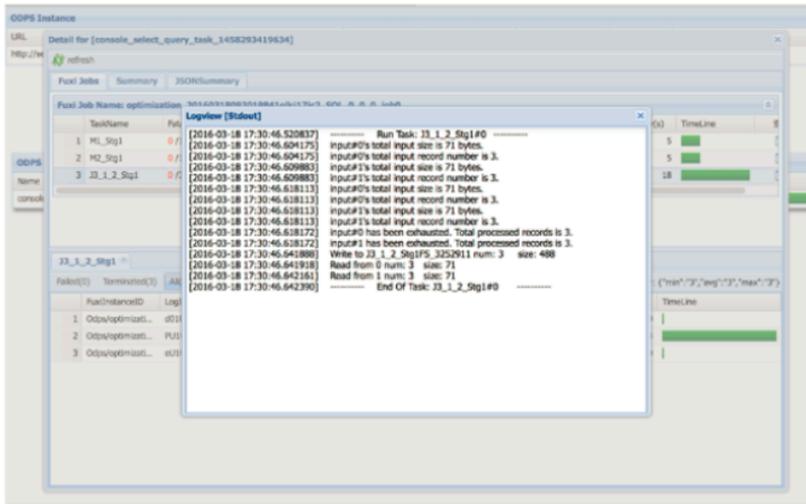


Different instances in the same Fuxi task should run for similar times. In this example, data skew occurs.

Click on stdout to see the amount of data processed, which can accurately determine the data skew, that is, the amount of data processed between different instances varies greatly.

View the UDF and MR debugging information

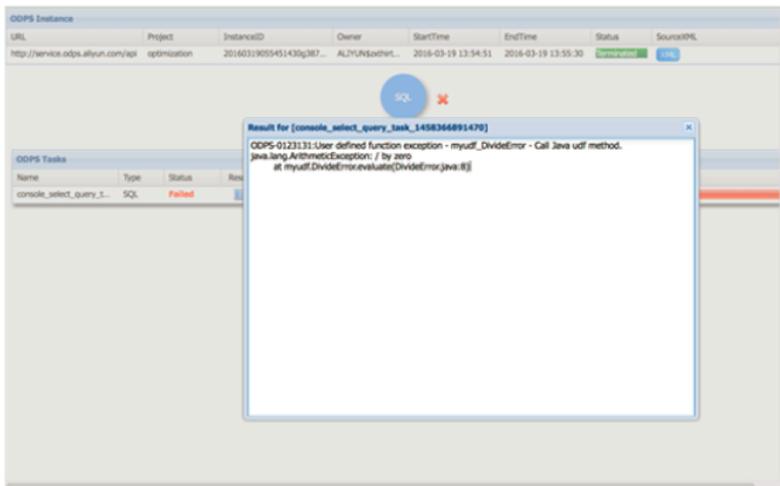
Figure 3-30: View the UDF and MR debugging information



View debugging information in Fuxi Instance Stuuout and Stderr

View the task status - Terminated

Figure 3-31: View the task status - Terminated



Error messages can be seen from the results of the job

You can also click Detail to go into details to see what went wrong.

3.2.2.4.5 Best practices

Locate LogView based on the instance ID

After you submit a job, you can press **Ctrl+C** to return to `odpscnd` and perform other operations. You can run the `wait <instanceid>`; command to locate LogView and obtain the job status.

Figure 3-32: Locate LogView based on the instance ID

```
odps@ optimization>select * from skew a join skew2 b on a.key=b.key;
ID = 20160318095028941gopbx6jc2
Log view:
http://logview.odps.aliyun.com/logview/?h=http://service.odps.aliyun.com/api&p=optimization&i=20160318095028941gopbx6jc2&token=U0ZBU1RwbGhmRES
jbnIN2gwY0lBMjFobjhrPSxPRFBTX09CTzoxMDE0DyNTI000IzNDUSLDE0NTg4OTk0NjkseyJTdGF0ZW1lbnQiOiI7Ikt7jdg1vbi16WyJvZHBz01JlYWQiXSwiRWZmZWNOIjo1QWxs
Bc1lCjSZXNvdXJzS16WyJhY3M6b2RwczoqOnByb2p1Y3RzL29wdGltaxphdGlvbi9pbmN0YNSjZXRhVWJAxNjAzNTgwTUwMjg5NDFn3BieDZqYzIiXX1dLCJWZXJzaW9uIjo1MSJ9
2016-03-18 17:50:40 M1_Stgl_job0:0/0/1[0%] M2_Stgl_job0:0/0/1[0%] J3_1_2_Stgl_job0:0/0/3[0%]
2016-03-18 17:50:45 M1_Stgl_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;
ID = 20160318095028941gopbx6jc2
Log view:
http://logview.odps.aliyun.com/logview/?h=http://service.odps.aliyun.com/api&p=optimization&i=20160318095028941gopbx6jc2&token=NVFFC1q2V1FSNmx
2TTNGeW9IL2QWU8z0UhfPSxPRFBTX09CTzoxMDE0DyNTI000IzNDUSLDE0NTg4OTk0NTcseyJTdGF0ZW1lbnQiOiI7Ikt7jdg1vbi16WyJvZHBz01JlYWQiXSwiRWZmZWNOIjo1QWxs
Bc1lCjSZXNvdXJzS16WyJhY3M6b2RwczoqOnByb2p1Y3RzL29wdGltaxphdGlvbi9pbmN0YNSjZXRhVWJAxNjAzNTgwTUwMjg5NDFn3BieDZqYzIiXX1dLCJWZXJzaW9uIjo1MSJ9
2016-03-18 17:50:58 M1_Stgl_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.
```

Locate running tasks

After you exit the control window, you can run the `show p`; command to locate currently running tasks and historical tasks.

Figure 3-33: Locate running tasks

StartTime	RunTime	Status	InstanceID	Owner	Query
2016-09-18 16:27:04	7s	Success	20160918082704275guto17jc2	ALIYUN\$	liyun.com select from dual;

3.2.2.5 Apsara Bigdata Manager

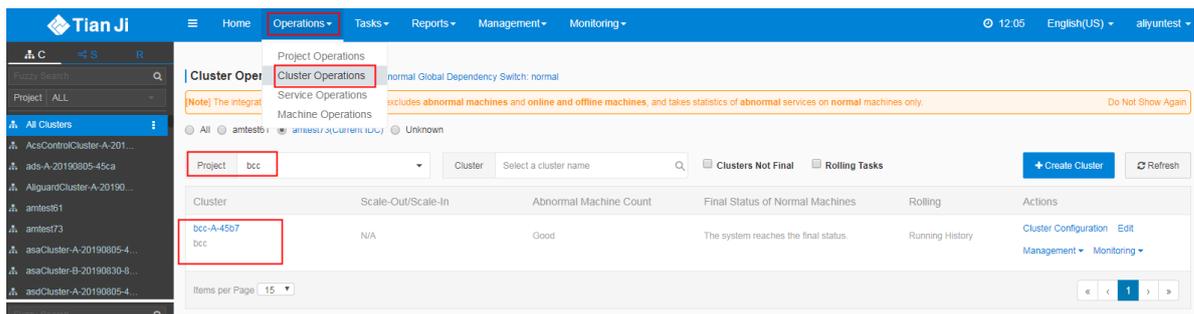
Apsara Bigdata Manager is an O&M management platform tailored to Alibaba big data services. Apsara Bigdata Manager currently supports services such as MaxCompute, Realtime Compute, DataWorks, AnalyticDB, and BigGraph.

Apsara Bigdata Manager provides service O&M functions in the form of service components. Each service component is made up of the service tree structure, configuration, automatic and manual service self-check, workflows, package management, global search, log search, indicator information, and metrics information, as well as some custom features and Apsara Stack Insight-specific features.

You can perform the following operations to log on to Apsara Bigdata Manager:

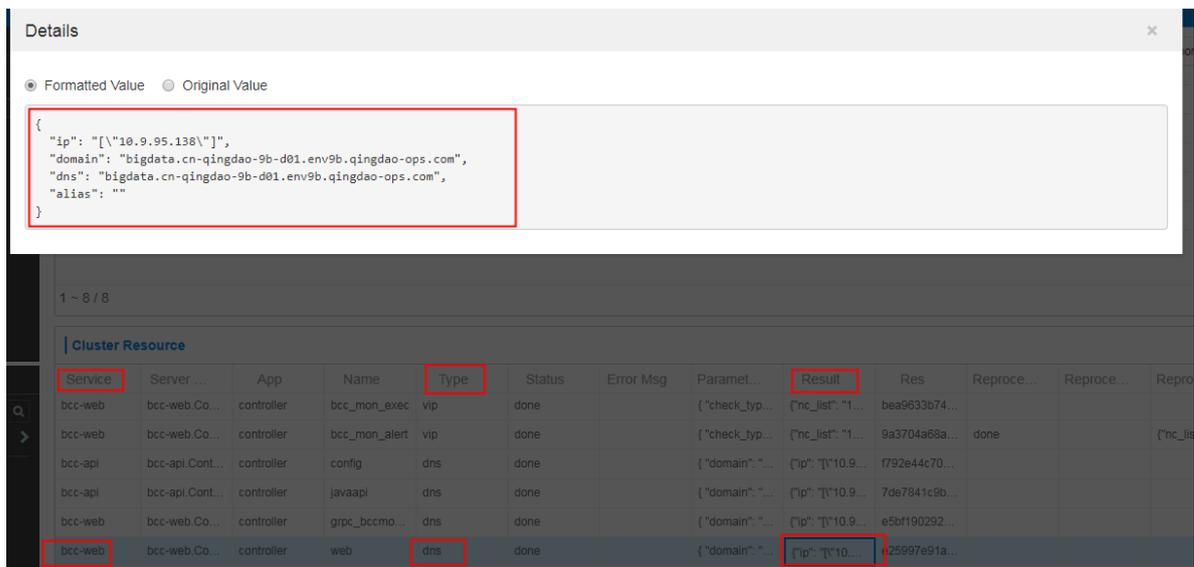
1. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. Select **bcc** from the Project drop-down list. Hover over the vertical dots next to the BCC cluster, and choose **Dashboard** from the shortcut menu. The Cluster Dashboard page is displayed.

Figure 3-34: Locate the BCC cluster



2. In **Cluster Resource**, set **Service** to **bcc-web** and **Type** to **dns** to filter the cluster resources. Right-click the **Result** column and choose **Show More** from the shortcut menu to obtain the domain name of Apsara Bigdata Manager.

Figure 3-35: Obtain the domain name of Apsara Bigdata Manager



3. Use the domain name you obtained from the preceding step to log on to Apsara Bigdata Manager.

Then the Apsara Bigdata Manager homepage is displayed.

For more information about how to use each feature available with Apsara Bigdata Manager, see *Apsara Bigdata Manager User Manual* in Help Center.

3.2.3 Routine O&M

3.2.3.1 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
messenger_partition_num=3
```

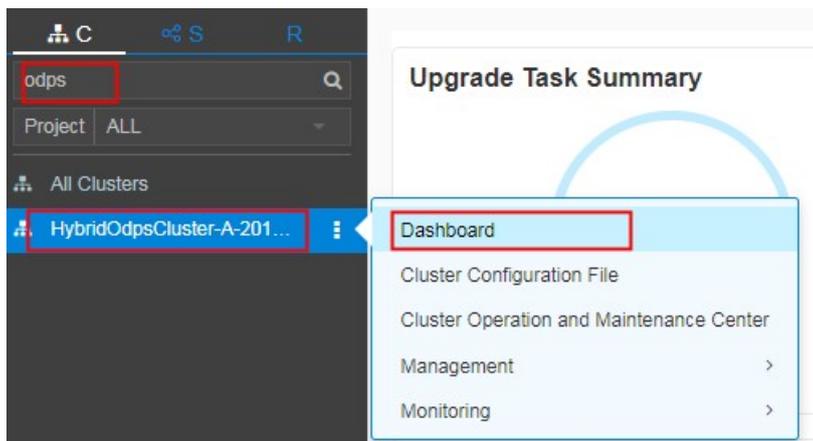
You can modify these parameter values based on your requirements and start the corresponding MaxCompute services based on the configured values. For more information, see *Restart a MaxCompute service*.

If you add `xstream_max_worker_num=3` at the end of the configuration file, XStream will be started with three running workers.

3.2.3.2 Routine inspections

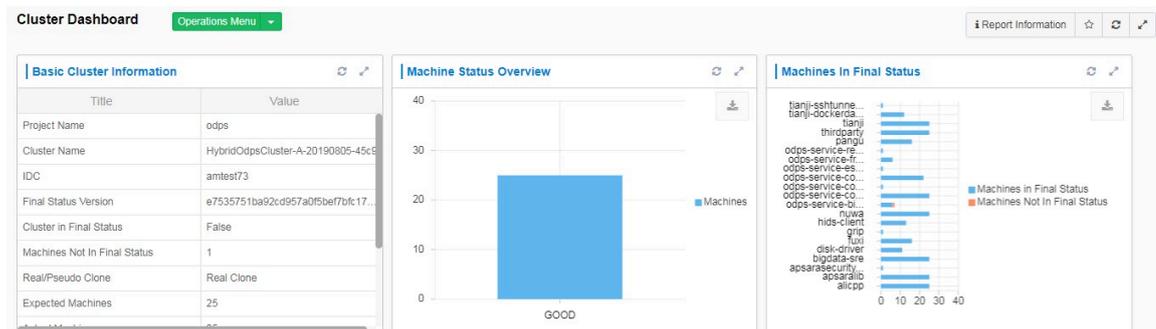
1. Check that all machines have reached the final state on the Cluster Operations and Maintenance Center page in Apsara Infrastructure Management Framework.
 - a. In the left-side navigation pane, click the C tab and search for odps. Hover over the vertical dots next to the MaxCompute cluster, and choose Dashboard from the shortcut menu.

Figure 3-36: Search for the odps cluster



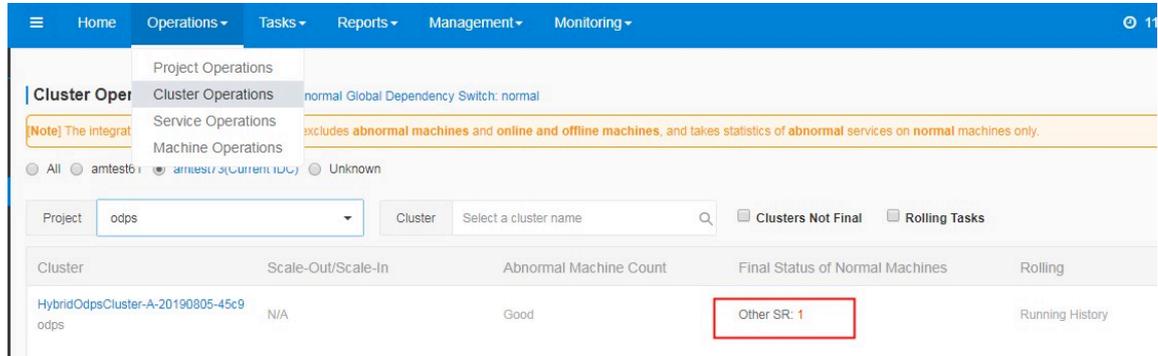
- b. On the Cluster Dashboard > Machines in Final Status page, check whether all machines have reached the final state. The following figure shows that some machines have not reached the final state.

Figure 3-37: Machines in Final State



- c. In the top navigation bar, choose **Operations > Cluster Operations**. Click **Other SR** to check the machines that have not reached the final state.

Figure 3-38: Other SR



- d. On the **Service Final Status Query** page, click **View Details** to view machine details.

Figure 3-39: Service Final Status Query

Cluster Operations > Service Final Status Query

odps-service-biggraph ^					
Server Role	Total Machines	Machines not in Final Status	Service Reason	Machine Reason	Actions
BigGraphFrontendServer#	3	0	0	0	View Details
BigGraphInit#	1	0	0	0	View Details
BigGraphWorker#	3	0	0	0	View Details
BigGraphZK#	3	0	0	0	View Details

odps-service-computer ^					
Server Role	Total Machines	Machines not in Final Status	Service Reason	Machine Reason	Actions
ComputerInit#	1	0	0	0	View Details
OdpsComputer#	6	0	0	0	View Details
PackageInit#	25	0	0	0	View Details
ServiceTest#	1	0	0	0	View Details

odps-service-console ^					
Server Role	Total Machines	Machines not in Final Status	Service Reason	Machine Reason	Actions
AdminConsole#	1	0	0	0	View 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;
ID = 20180420061754827g78x7i
Log view:
http://logview.cn-hangzhou-env6-d01.odps.aliyun-inc.com:9000/logview/?h=http://s
180420061754827g78x7i&token=aEVmNTF1dm5GMnFOV1BSWjViZE0rOWRERnZFPSxPRFBTX09CTzox
SwiRWZmZWN0IjoiQWxsY3ciLCJSZXNvdXJzSI6WyJhY3M6b2RwczoqOnByb2p1Y3RzL29kcHNfc21va
J9
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      |
-----+
1        |
```

As shown in the following figure, `fluxi job` is running. The command output indicates that the cluster is functioning properly.

```

odps@ odps_smoke_test> select count(*) from datahub_smoke_test
>;

ID = 20180420065305115gv5pf9d
Log view:
http://logview.cn-beijing-bgm-d01.odps.bgm.com:9000/logview/?h=http://servic
80420065305115gv5pf9d&token=VS9hRzc4RjAzeXJ2bmRF0utyYnNWSXfKw0wPSxPRFBTX090
iI6WyJvZHBz0lJlYwQiXSwiRWZmZWN0IjoiQWxsY3ciLCJSZXNvdXJjZSI6WyJhY3M6b2Rwcz0q
UzMDUxMTVndjVwZjlkIl19XSwiVmVyc2lvbiI6IjEifQ==
2018-04-20 14:53:10 M1_Stgl_job0:0/0/1[0%]      R2_1_Stgl_job0:0/0/1[0%]
2018-04-20 14:53:15 M1_Stgl_job0:0/1/1[100%]   R2_1_Stgl_job0:0/0/1[0%]
2018-04-20 14:53:20 M1_Stgl_job0:0/1/1[100%]   R2_1_Stgl_job0:0/1/1[100%]
2018-04-20 14:53:25 M1_Stgl_job0:0/1/1[100%]   R2_1_Stgl_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
M1_Stgl:
  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_Stgl: (min: 0, max: 0, avg: 0)
R2_1_Stgl:
  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 if they have been restarted recently:

a. `r swl Odps/MessengerServicex`

```
$r swl Odps/MessengerServicex
WorkerName | LastUpdateTime | pid | planned | loaded | unloaded
MessageServerRole@101h05215.cloud.h07.amtest1284 | Mon Apr 9 16:49:03 2018 | 24697 | 1 | 1 | 0
MessageServerRole@101h11210.cloud.h13.amtest1284 | Mon Apr 9 16:48:37 2018 | 15149 | 1 | 1 | 0
MessageServerRole@101h08109.cloud.h09.amtest1284 | Mon Apr 9 16:49:03 2018 | 23586 | 1 | 1 | 0
```

b. `r swl Odps/OdpsServicex`

```
$r swl Odps/OdpsServicex
WorkerName | LastUpdateTime | pid | planned | loaded | unloaded
RecycleWorker@101h08114.cloud.h09.amtest1284 | Mon Apr 9 17:05:42 2018 | 52905 | 0 | 0 | 0
OdpsWorker@101h08114.cloud.h09.amtest1284 | Mon Apr 9 17:05:42 2018 | 52904 | 0 | 0 | 0
OdpsWorker@101h11010.cloud.h11.amtest1284 | Mon Apr 9 17:04:06 2018 | 4454 | 0 | 0 | 0
ExecutorWorker@101h08114.cloud.h09.amtest1284 | Mon Apr 9 17:05:42 2018 | 52903 | 0 | 0 | 0
ExecutorWorker@101h11010.cloud.h11.amtest1284 | Mon Apr 9 17:04:22 2018 | 6524 | 0 | 0 | 0
SchedulerWorker@101h08114.cloud.h09.amtest1284 | Mon Apr 9 17:05:47 2018 | 53609 | 0 | 0 | 0
WorkflowWorker@101h08114.cloud.h09.amtest1284 | Mon Apr 9 17:05:48 2018 | 53610 | 0 | 0 | 0
```

c. `r swl Odps/HiveServerx`

```
$r swl Odps/HiveServerx
WorkerName | LastUpdateTime | pid | planned | loaded | unloaded
AuthServer@101h08114.cloud.h09.amtest1284 | Tue Apr 10 18:05:54 2018 | 23585 | 0 | 0 | 0
HiveServer@101h11010.cloud.h11.amtest1284 | Mon Apr 9 17:03:07 2018 | 1696 | 1 | 1 | 0
HiveServer@101h08114.cloud.h09.amtest1284 | Tue Apr 10 18:06:02 2018 | 23587 | 2 | 2 | 0
CatalogServer@101h08114.cloud.h09.amtest1284 | Tue Apr 10 18:05:55 2018 | 23586 | 1 | 1 | 0
```

d. `r swl Odps/QuotaServicex`

```
$r swl Odps/QuotaServicex
WorkerName | LastUpdateTime | pid | planned | loaded | unloaded
QuotaWorkerRole@101h08114.cloud.h09.amtest1284 | Mon Apr 9 16:55:32 2018 | 32814 | 0 | 0 | 0
```

e. `r swl Odps/ReplicationServicex`

```
$r swl Odps/ReplicationServicex
WorkerName | LastUpdateTime | pid | planned | loaded | unloaded
ReplicationServer@101h05215.cloud.h07.amtest1284 | Mon Apr 9 16:49:12 2018 | 26594 | 0 | 0 | 0
ReplicationServer@101h11210.cloud.h13.amtest1284 | Mon Apr 9 16:48:51 2018 | 26859 | 0 | 0 | 0
ReplicationServer@101h11215.cloud.h13.amtest1284 | Mon Apr 9 16:49:18 2018 | 3453 | 0 | 0 | 0
ReplicationMaster@101h11010.cloud.h11.amtest1284 | Mon Apr 9 16:50:21 2018 | 34315 | 0 | 0 | 0
```

4. Run the following command to check whether any errors exist:

```
puadmin lscs |grep -vi NORMAL|grep -vi DISK_OK
```

```
puadmin lscs |grep -vi NORMAL|grep -vi DISK_OK
The pangou disk status:
Total Disk Size:681225 GB
Used Free Disk Size:635009 GB
Total File Size:1093 GB
Total UnReserved Disk Space4Piops:0 GB
Total Disk Space4Piops:0 GB
Total UnReserved Disk Tops4Piops:0
Total Disk Tops4Piops:0
TotalChunkNumber:26074944 NonTempChunkNumber:26074030 NonTempChunkDataSize:1093 GB TempChunkNumber:914 TempChunkDataSize:0 GB
No. Rack UsableChunkserver/TotalChunkserver UsableDisk/TotalDisk TotalDiskSize TotalFreeDiskSize
1 101g15 2/2 23/23 128427 GB 119672 GB
2 101h05 1/1 11/11 61421 GB 57318 GB
3 101h09 2/2 23/23 150763 GB 140758 GB
4 101h11 5/5 57/57 340612 GB 317859 GB
Number of Racks: 4
Number of Usable Racks(Having at least one disk with Free Disk Size > 15GB): 4
Notice!: Total Disk Size of 101h11 >= 1/3 of Total Disk Size of the Cluster, three replicas may not locate in different racks
```

5. Run the following commands to check for data integrity:

- a. `puadmin fs -abnchunk -t none`

```
$puadmin fs -abnchunk -t none
Master Address: nuwa://localcluster/sys/pangu/master
ChunkId Type      FoundTime
```

- b. `puadmin fs -abnchunk -t onecopy`

```
$puadmin fs -abnchunk -t onecopy
Master Address: nuwa://localcluster/sys/pangu/master
ChunkId Type      FoundTime
```

- c. `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"
```

```
$tj_show -r nuwa.NuwaZK#>/tmp/nuwa;pssh -h /tmp/nuwa -i "echo srvr | nc localhost 10240 | grep Mode"
[1] 15:59:01 [SUCCESS] vm010036016093
Mode: follower
[2] 15:59:02 [SUCCESS] vm010036032042
Mode: leader
[3] 15:59:02 [SUCCESS] vm010036024022
Mode: follower
```

7. Run the following commands to check that Apsara Distributed File System functions properly:

```
puadmin gems
```

```
puadmin gss
```

```

puadmin gems
ElectMasterStatus : ELECT_MASTER_OVER_ELECTION
PrimaryId         : tcp://10.36.8.33:10260
PreferedWorkerid  :
PrimaryLogId      : 617851602
TotalWokerNumber  : 3
ElectConsentNumber : 2
SyncConsentNumber : 2
ElectSequence     : [935155f0-fb68-4cd9-bee9-08d23afe84eb,4,1328760004]
WorkerStatus      :
  tcp://10.36.16.92:10260 : ELECT_WORKER_STATUS_SECONDARY
  tcp://10.36.32.41:10260 : ELECT_WORKER_STATUS_SECONDARY
  tcp://10.36.8.33:10260  : ELECT_WORKER_STATUS_PRIMARY

[admin@sm010036032037 /home/admin]
puadmin gss
PrimaryStatus : PRIMARY_STARTUP_SERVICE_STARTED
PrimaryCurrentLogId : 617852679
WorkerSyncStatus :
  tcp://10.36.16.92:10260[SyncedLogId:617852670, LastFailTime:2018-04-17 12:07:43, WorkerType: NORMAL]
  tcp://10.36.32.41:10260[SyncedLogId:617852638, LastFailTime:1970-01-01 08:00:00, WorkerType: NORMAL]

```

8. Perform daily inspections in Apsara BigData Manager to check disk usage.

3.2.3.3 Chunkserver and tunnel scale-in

Potential risks

Cluster resources are reduced as a result of the scale-in. This operation may affect the originally configured quota. You must evaluate the original quota configurations after the scale-in. If there is only the default quota, you do not have to evaluate the configurations. You only need to inform customers about the amount of remaining resources.

Procedure

1. Perform the on-site check.

- a. In Apsara Infrastructure Management Framework, locate Adminconsole# in the odps-service-console service of the odps cluster, and open the corresponding TerminalService window.
- b. Run the `r ttrll` command to confirm the role information of the machine to be removed.

Figure 3-40: Confirm the role information of the machine to be removed

```

$ r ttrll
total tubo in cluster=11

detail table for every machine:
Machine Name      | CPU   | Memory | Other
-----
                    | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
                    | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
                    | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
                    | 6,300 | 170,453 | ElasticSearchInstance:5
                    | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
                    | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
                    | 6,300 | 170,453 | OdpsSpecialInstance:20 OdpsCommonInstance:20
                    | 6,300 | 170,453 | ElasticSearchInstance:5
                    | 6,300 | 170,453 | ElasticSearchInstance:5
                    | 6,300 | 170,453 | OdpsSpecialInstance:20 OdpsCommonInstance:20
                    | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
Total              | 69,300 | 1,874,983 | NA
    
```

- c. Log on to the machine to be removed and run the `tj_show |grep -iv tunnel` command to confirm that no tunnel information is available on the machine.

Figure 3-41: Confirm that no tunnel information is available on the machine

```

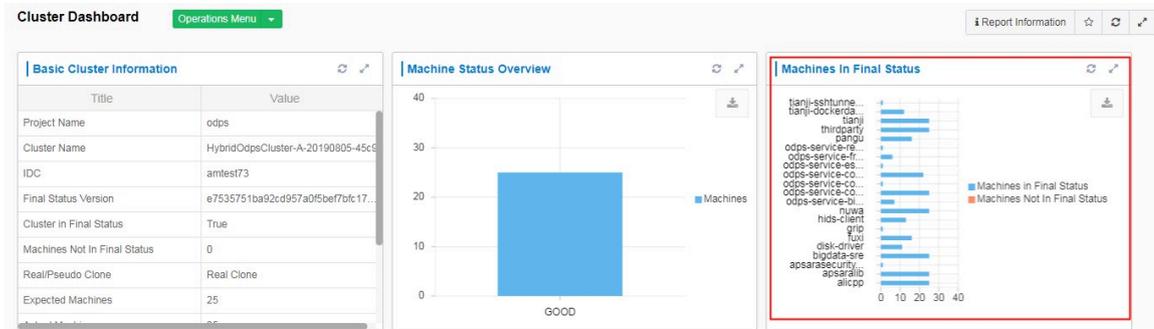
$tj show |grep -iv tunnel
Local_hostname:
Local_address:
Local_cluster:
Local_serverroles: apsaralib.ApsaraLib#016|apsaralib.ApsaraLib#legacy|bigdata-sre.Agent#|disk-driver
.DiskDriverWorker#|fuxi.DeployAgent#|fuxi.FuxiMonitor#|fuxi.Tubo#|hids-client.HidsClient#|nuwa.NuwaC
onfig#|odps-service-computer.OdpsComputer#|odps-service-computer.PackageInit#|odps-service-controlle
r.checkComputer#|odps-service-controller.checkCpuStatus#|pangu.PanguChunkserver#|pangu.PanguMonitor#
|thirdparty.ThirdpartyLib#|tianji.TianjiClient#
    
```

2. Before you scale in a cluster, make sure that the cluster has reached its final state and is functioning properly.

- a. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose Operations > Cluster Operations. On the Cluster

Operations page, click the cluster you want to scale in to go to the cluster dashboard. Check for machines that have not reached the final state.

Figure 3-42: Check for machines that have not reached the final state



b. Run the following commands to check whether the data and replicas in Apsara Distributed File System are secure:

```
/apsara/deploy/puadmin abnchunk
```

```
/apsara/deploy/puadmin fs -abnchunk -t lessmin
```

```
puadmin fs -abnchunk -t none
```

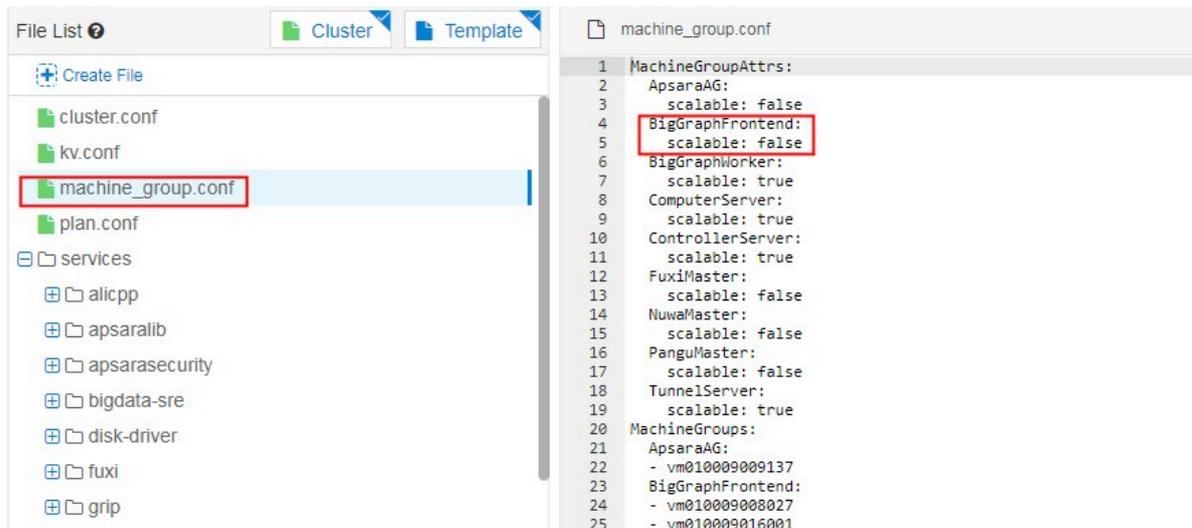
```
puadmin fs -abnchunk -t onecopy
```

```
puadmin gems
```

```
puadmin gss
```

3. Make sure that the scalable tag value is true for the cluster to be scaled in. If the value is false, manually change the value to true in the cluster configuration file and submit a rolling task.

Figure 3-43: Confirm the scalable tag value



4. Run the following command to set the machine status to shutdown in Apsara Distributed File System:

```
puadmin cs -stat tcp://ip:10260 --set=shutdown
```

5. Run the following commands on the OPS1 server to scale in the service cluster:

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
./tianji_ops_tool.py contract_nc -c clusterName -l machineList --
config tianji_clt.conf -s SRGname
```

-- You must run the commands as the root user. If you run the commands as an admin user, an error occurs.

**Note:****Parameters:**

- **-c:** required. The name of the cluster to remove the servers from.
- **-s:** required. The name of the server role group (SRG) to remove the servers from.
- **--config:** required. The `tianji_clt` configuration file.
- **-l:** required. The hostnames of the servers to be removed. Separate multiple hostnames with commas (,).

Command output:

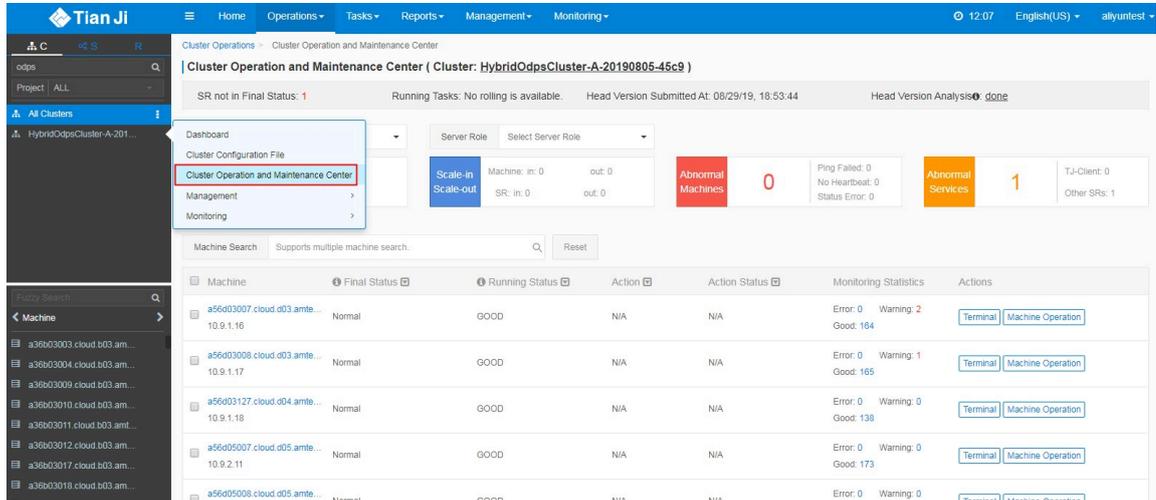
```
machines removed from cluster XXX, newversion: 53111bf22f40e64d5343
e0a73a2bce3241b7dac5
-- Indicates that the operation is completed.
```

6. Check whether the operation has taken effect in Apsara Infrastructure Management Framework.

- In the top navigation bar, choose Operations > Cluster Operations. On the upper part of the left-side navigation pane, hover over the vertical dots next to**

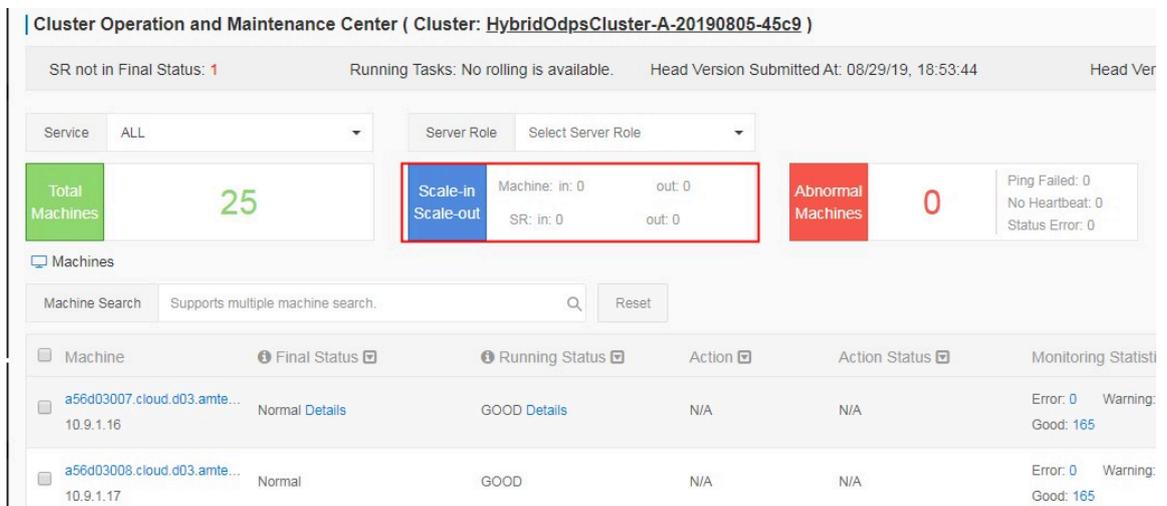
the target cluster and choose Cluster Operation and Maintenance Center from the shortcut menu.

Figure 3-44: Cluster Operation and Maintenance Center



b. Check whether the machine has been removed.

Figure 3-45: Check whether the machine has been removed



Note:

The preceding steps are used to remove a machine from the service cluster and add it to the default Apsara Infrastructure Management Framework cluster. If you need to remove the machine from Apsara Infrastructure Management Framework, perform the following steps.

7. Remove the machine from Apsara Infrastructure Management Framework.

- a. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Machine Operations**. On the Machine Operations page that appears, click **Machine Online/Offline**.
- b. On the Machine Online/Offline page that appears, click **Remove Machine**. In the Enter Machine List area on the left, enter the hostname of the machine. If the machine can be removed, click **Clear Machines**.

8. After you complete the scale-in, run the following command to modify the quota:

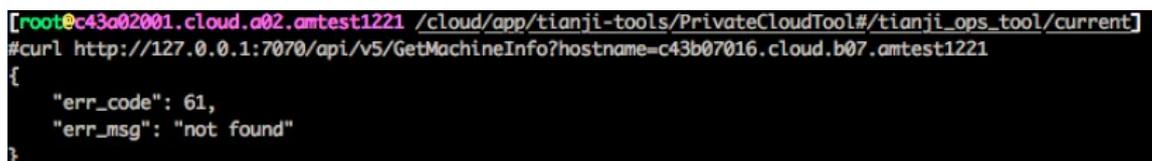
```
r setquota -i ${Account} -a ${Alias} -s ${CPU_Quota_New} ${MEM_Quota_New}
```

9. Verify the result.

- Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Machine Operations**. Search for the machine hostname and check whether this machine belongs to the default cluster. If yes, the machine has been removed from the service cluster to the default cluster.
- Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. Enter the machine hostname to check whether the machine exists. If no data is found, the machine has been removed from Apsara Infrastructure Management Framework.
- Run the following API commands on the OPS1 server to check whether the machine has been removed from Apsara Infrastructure Management Framework:

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
curl http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=$hostname
```

Figure 3-46: Verify the result by using API commands



```
[root@c43a02001.cloud.a02.amtest1221 /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current]
#curl http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=c43b07016.cloud.b07.amtest1221
{
  "err_code": 61,
  "err_msg": "not found"
}
```

10. Verify the existence of the removed machine.

- a. Run the `r ttrl` command in the `TerminalService` window to confirm that the machine does not exist.
- b. Run the following commands to check whether the data and replicas in Apsara Distributed File System are secure:

```
/apsara/deploy/puadmin abnchunk
```

```
/apsara/deploy/puadmin fs -abnchunk -t lessmin
```

```
puadmin fs -abnchunk -t none
```

```
puadmin fs -abnchunk -t onecopy
```

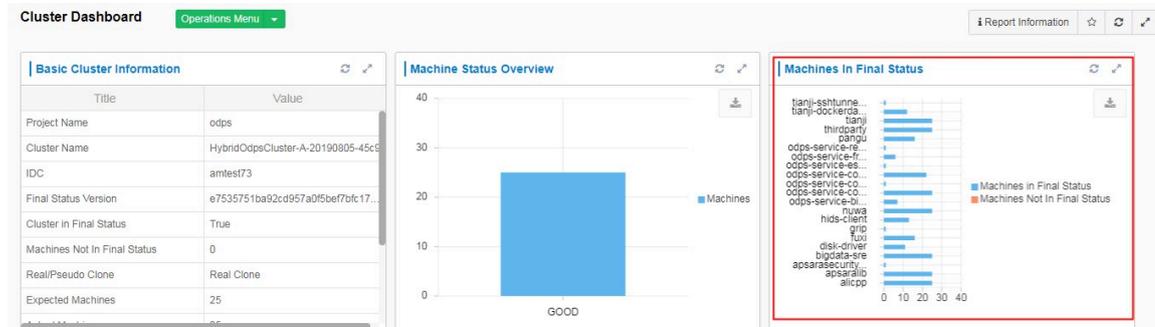
3.2.3.4 Chunkserver and tunnel scale-out

Procedure

1. In Apsara Infrastructure Management Framework, locate `Adminconsole#` in the `odps-service-console` service of the `odps` cluster, and open the corresponding `TerminalService` window.
2. Before you perform the scale-out, make sure that the cluster has reached the final state and is functioning properly.
 - a. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose `Operations > Cluster Operations`. On the Cluster

Operations page, click the cluster you want to scale out to go to the cluster dashboard. Check for machines that have not reached the final state.

Figure 3-47: Check for machines that have not reached the final state



- b. Run the following commands to check whether the data and replicas in Apsara Distributed File System are secure:**

```
/apsara/deploy/puadmin abnchunk
```

```
/apsara/deploy/puadmin fs -abnchunk -t lessmin
```

```
puadmin fs -abnchunk -t none
```

```
puadmin fs -abnchunk -t onecopy
```

```
puadmin gems
```

```
puadmin gss
```

- When you add a buffer cluster, make sure that the new IP addresses that Deployment Planner assigns to the machines are not used in the current environment. This can avoid exceptions arising from IP address conflicts after the scale-out.
- Because clone protection is configured in Apsara Infrastructure Management Framework, you must set `clone_mode` to `normal` before you scale out a cluster, and set `clone_mode` to `block` after the scale-out is completed. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations > Global Clone Switch**.
- Add the buffer cluster.



Note:

When you make plans for the scale-out in Deployment Planner, make sure that the new buffer cluster name is unique.

- a. Copy `_tianji_imports` to the `/apsarapangu/disk3/u_disk/` directory of the OPS1 server. Run the following command in the `tianji_zhuque_sdk` directory:

```
./tianji_zhuque_exchanger.py import --skip_packages -o ${final
status in Apsara Infrastructure Management Framework} -c
tianji_dest.conf
```

- b. Log on to the Apsara Infrastructure Management Framework console. In the left-side navigation pane, locate the buffer cluster in the cluster list. Then, hover over the vertical dots next to the buffer cluster, and choose **Cluster Operations and Maintenance Center** from the shortcut menu to view the status of machines.

Figure 3-48: View the status of machines

The screenshot shows the 'Cluster Operation and Maintenance Center' for a cluster named 'HybridOdpsCluster-A-20190805-45c9'. Key metrics include:

- Total Machines: 25
- Abnormal Machines: 0
- Abnormal Services: 1
- Machine In/Out: 0
- SR In/Out: 0
- Ping Failed: 0, No Heartbeat: 0, Status Error: 0
- T-J-Client: 0, Other SRs: 1

A table lists the status of individual machines:

Machine	Final Status	Running Status	Action	Action Status	Monitoring Statistics	Actions
a56d03007.cloud.d03.amte... 10.9.1.16	Normal	GOOD	N/A	N/A	Error: 0, Warning: 0, Good: 166	Terminal, Machine Operation
a56d03008.cloud.d03.amte... 10.9.1.17	Normal	GOOD	N/A	N/A	Error: 0, Warning: 0, Good: 166	Terminal, Machine Operation
a56d03127.cloud.d04.amte... 10.9.1.18	Normal	GOOD	N/A	N/A	Error: 0, Warning: 0, Good: 138	Terminal, Machine Operation
a56d05007.cloud.d05.amte... 10.9.2.11	Normal	GOOD	N/A	N/A	Error: 0, Warning: 2, Good: 171	Terminal, Machine Operation

- c. Run the following API commands to view information about the scale-out:

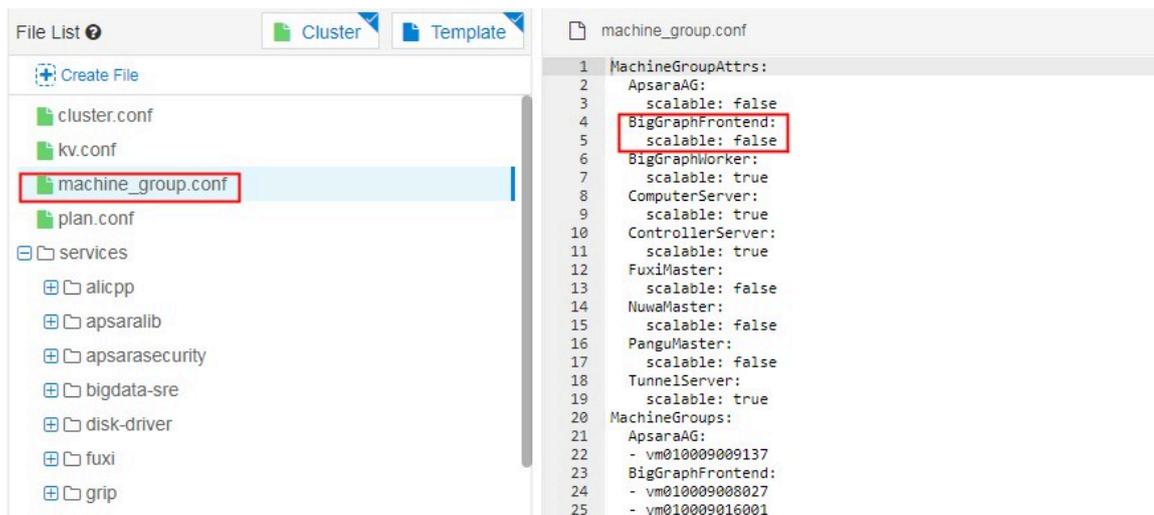
```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/
current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
```

```
./tianji_clt machinestatus -c buffer --config clt2.conf
```

6. Perform a scale-in by removing a machine from the service cluster to the default Apsara Infrastructure Management Framework cluster.

- a. Make sure that the scalable tag value is true for the new buffer cluster. If the value is false, manually change the value to true in the cluster configuration file and submit a rolling task.**

Figure 3-49: Confirm the scalable tag value



- b. Run the following commands on the OPS1 server to scale in the service cluster:**

```

cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/
current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
./tianji_ops_tool.py contract_nc -c clusterName -l machine1,
machine2 --config clt2.conf -s SRGname
-- You must run the commands as the root user. If you run the
commands as an admin user, an error occurs.

```



Note:

Parameters:

- **-c:** required. The name of the cluster that the servers are removed from.
- **-s:** required. The name of the server role group (SRG) that the servers are removed from.
- **--config:** required. The tianji_clt configuration file.

- **-l: required. The hostnames of the servers to be removed. Separate multiple hostnames with commas (,).**

The command output is as follows:

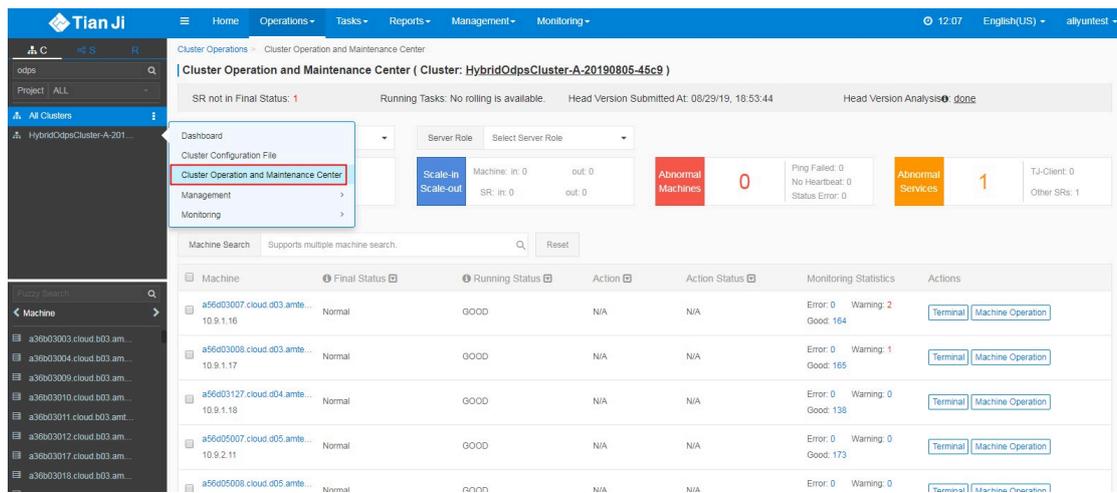
```
machines removed from cluster XXX, newversion: 53111bf22f  
40e64d5343e0a73a2bce3241b7dac5
```

-- Indicates that the operation is completed.

c. Check whether the operation takes effect in Apsara Infrastructure Management Framework.

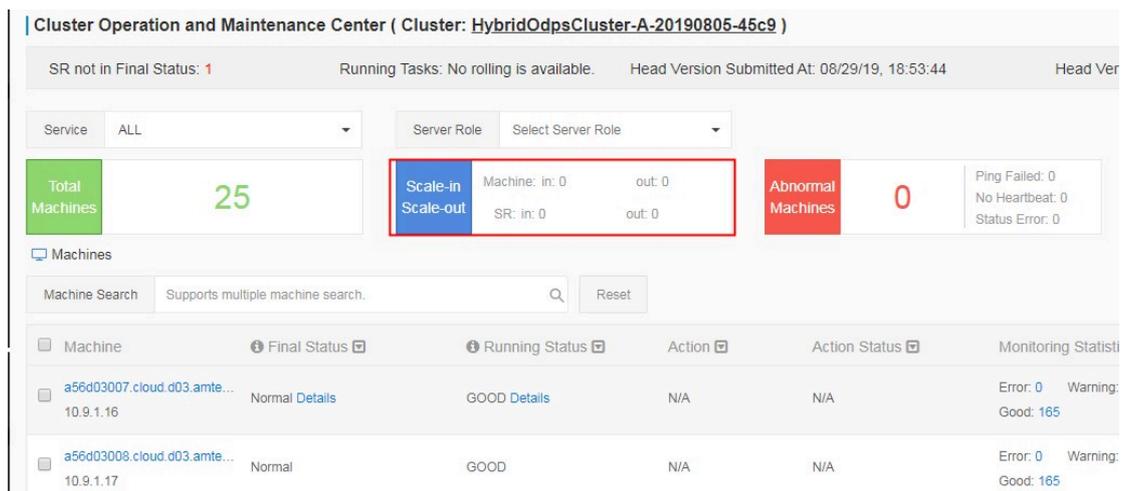
A. In the top navigation bar, select Operations > Cluster Operations. On the upper part of the left-side navigation pane, hover over the vertical dots next to the target cluster and click Cluster Operation and Maintenance Center.

Figure 3-50: Cluster Operation and Maintenance Center



B. Check whether the servers have been removed from the cluster.

Figure 3-51: Cluster Operation and Maintenance Center



d. Run the following commands to check the server removal task.

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/
current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
```

```
./tianji_clt machinestatus -c default --config clt2.conf
```

- e. Access the Cluster Configuration page, and check whether the machine in the cluster has been deleted from the `machine_group.conf` configuration file. If the machine still exists in the file, manually delete the machine and submit a rolling task.
 - f. Delete the new buffer cluster.
7. Add the machine to the XXX cluster, and specify the SRG to which the machine belongs.
- a. Check whether the clone mode for the XXX cluster is Real Clone.

Figure 3-52: Confirm the clone mode

Basic Cluster Information	
Title	Value
Project Name	odps
Cluster Name	HybridOdpsCluster-A-20190805-45c9
IDC	amtest73
Final Status Version	e7535751ba92cd957a0f5bef7bfc17..
Cluster in Final Status	True
Machines Not In Final Status	0
Real/Pseudo Clone	Real Clone
Expected Machines	25

- b. Run the following commands to perform the scale-out operation. The rolling task will be triggered after the commands are executed. The following command output indicates that the base cluster has been scaled out.

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/
current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
./tianji_ops_tool.py expand_nc -c XXXCluster-A-0000 -s SRG -l
machine1,machine2 --config clt2.conf
-- You must run the commands as the root user. If you run the
commands as an admin user, an error occurs.
```

**Note:****Parameters:**

- **-c: required.** The name of the cluster to be scaled out.

- **-s:** required. The name of the SRG that exists in Apsara Infrastructure Management Framework and can be found in the `machine_group.conf` file. If the SRG does not exist, you must use Deployment Planner to deploy the SRG.
- **-l:** required. The hostname of the machine to be added to the cluster. Separate multiple hostnames with commas (,).

Figure 3-53: Command output for successful scale-out

```
[root@0a1a03201.cloud.a05.am54 /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current] JG4
#ln -s /cloud/data/bootstrap_controller/BootstrapController#/bootstrap_controller/tianji_dest.conf clt2.conf

[root@0a1a03201.cloud.a05.am54 /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current] JG4
#./tianji_ops_tool.py expand_nc -C BasicCluster-20171011-b61f -s BaseC4pGwGroup -l 0a1a06205.cloud.a08.am54 --config clt2.conf
machines extended into cluster BasicCluster-20171011-b61f, newversion: 592ec2e75aacc69783eea7fb9d0848d7594c941f success
```

- c. Run the following command to check the cluster to which the machine belongs and the uplink information:

```
curl http://127.0.0.1:7070/api/v3/column/m.*?m.id=[machine hostname]
```

- d. Log on to the OdpsClone 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 -
```

- e. Check the task status in the operation log. If the status becomes rolling succeeded, the scale-out is completed.
8. After you complete the scale-out, export the latest final state from Apsara Infrastructure Management Framework to Deployment Planner to ensure the success of subsequent scale-in and scale-out operations.

9. After you complete the scale-out, run the following command to modify the quota:

```
r setquota -i ${Account} -a ${Alias} -s ${CPU_Quota_New} ${MEM_Quota_New}
```

10. Verify the result.

- Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose Operations > Cluster Operations. Check whether the cluster has reached the final state.
- Check whether the machine has been added to the SR to which the machine belongs.

Figure 3-54: Check whether the server has been applied with the specified server role



11. Check the monitoring data of the newly added server.

Log on to the Apsara Infrastructure Management Framework console. Navigate to the Server Details page of the newly added server and check whether there is monitoring data. If no monitoring data is available, restart the following three SRs: AcceleratorAgg#, AcceleratorMaster#, and AcceleratorSource#.

12. Verify the existence of the newly added machine.

Run the `r ttrtl` command in the TerminalService window to check that the newly added machine exists.

Rollback plan

If the physical machine in ComputerServer fails to be added to the expected cluster, remove the physical machine to the default cluster. The procedure is as follows:



Notice:

Make sure that the physical machine removed to the default cluster is the one that fails to be added to the expected cluster.

1. Remove the physical machine that fails to be added to the expected cluster.

On the ops1 server, run the following commands to remove a server to the default cluster. If you need to remove multiple servers to the default cluster, run the commands to remove them one by one.

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
./tianji_ops_tool.py contract_nc -c HybridOdpsCluster-A-xxxx -s
ComputerServer -l machine1 --config clt2.conf
./tianji_ops_tool.py contract_nc -c HybridOdpsCluster-A-xxxx -s
ComputerServer -l machine2 --config clt2.conf
```



Note:

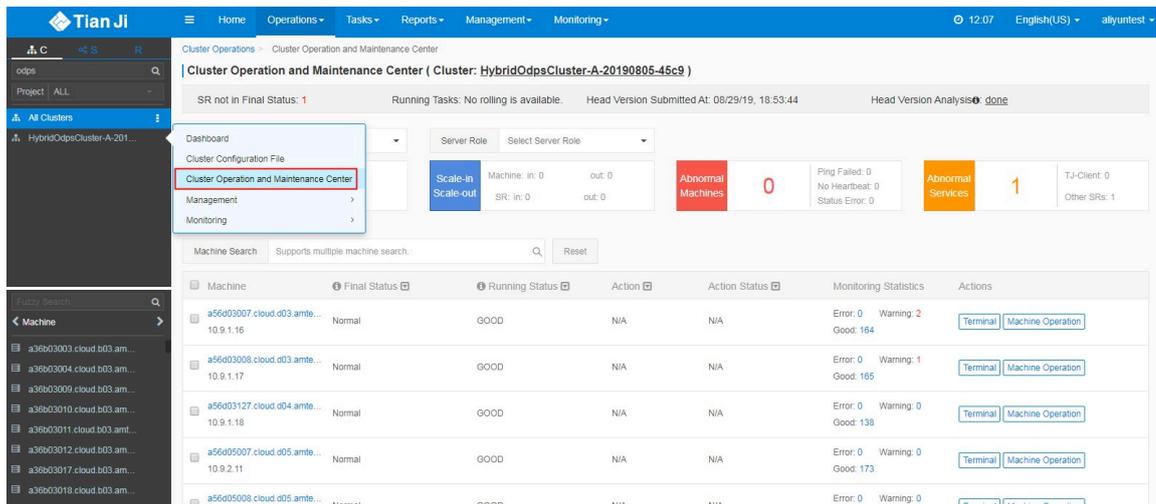
Parameters:

- **-c:** required. The name of the cluster that servers are removed from.
- **-s:** required. The name of the SRG that servers are removed from.
- **--config:** required. The tianji_clt configuration file.
- **-l:** required. The hostname of the server that failed to be added to the specified cluster.

2. Check whether the rollback operation takes effect.

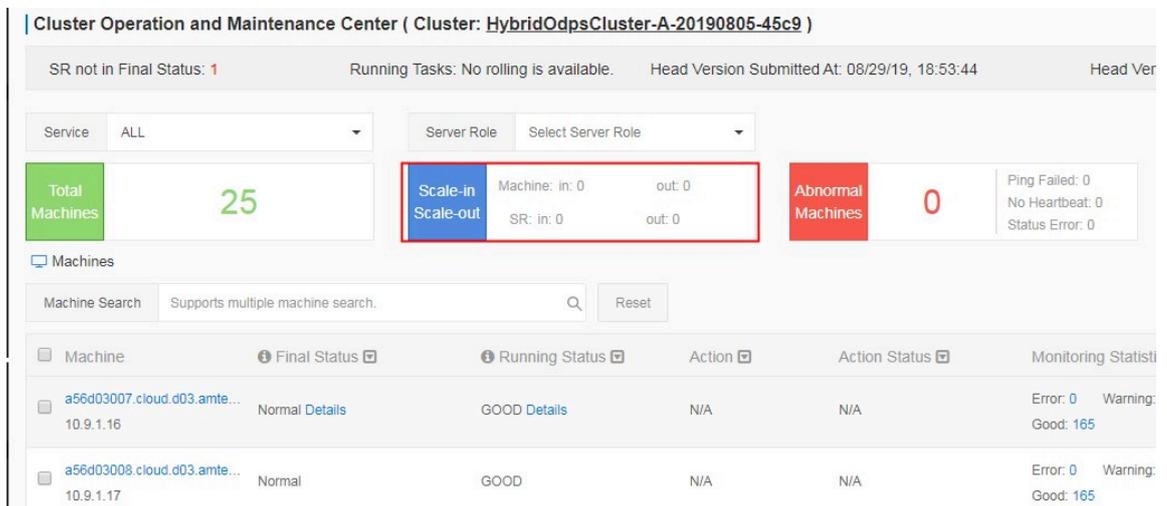
- a. In the top navigation bar, select Operations > Cluster Operations. On the upper part of the left-side navigation pane, hover over the vertical dots next to the target cluster and click Cluster Operation and Maintenance Center.

Figure 3-55: Cluster Operation and Maintenance Center



- b. Check whether the servers have been removed from the cluster.

Figure 3-56: Cluster Operation and Maintenance Center



3. Run the following commands to check the server removal task.

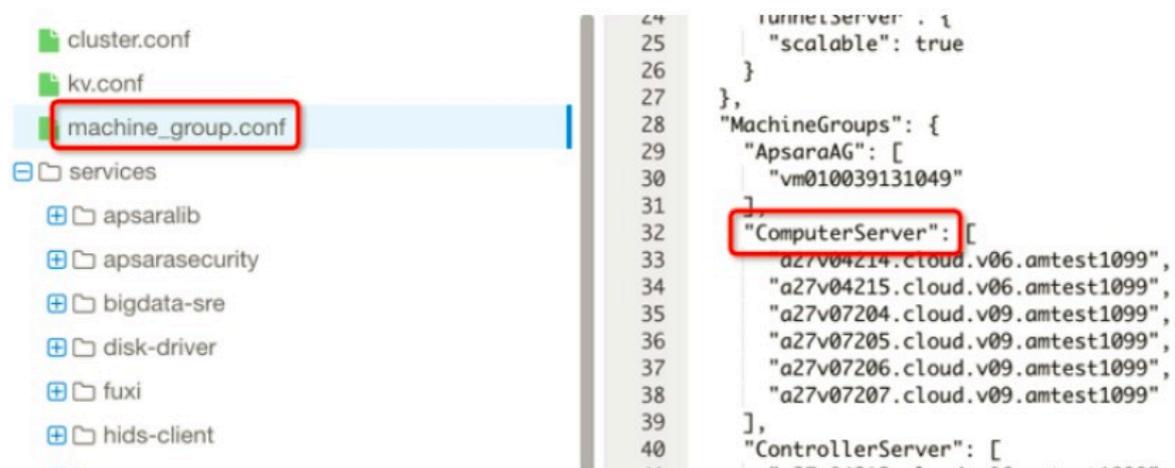
```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
```

```
./tianji_clt machinestatus -c default --config clt2.conf
```

4. Check whether the physical machine information has been deleted from the cluster configuration.

Navigate to the cluster configuration page of the XXX cluster, and check whether the physical machine information has been deleted from the cluster configuration. If any machine information is still in the machine_group.conf file, you must manually delete the information, submit a rolling task, and wait for the rolling process to complete.

Figure 3-57: Check whether the physical machine information has been deleted from the cluster configuration



5. If the quota has been updated, run the following command to roll the quota back to the value before the scale-out:

```
r setquota -i ${Account} -a ${Alias} -s ${CPU_Quota_Old} ${MEM_Quota_Old}
```

3.2.3.5 Shut down a chunkserver, perform maintenance, and then clone the chunkserver

Prerequisites

- A customer asks to fix faulty odps_cs and clone a new one.
- You must inform the customer that this operation will make one chunkserver in a cluster temporarily unavailable, but will not affect the service.
- All MaxCompute services have reached the final state and are functioning properly.

- All services on the OPS1 server have reached the final state and are functioning properly.
- You must ensure that the disk space available is sufficient for data migration triggered when a node goes offline.
- Check whether the primary node exists on the machine to be brought offline. If yes, switch services from the primary node to the secondary node.

Procedure

1. In Apsara Infrastructure Management Framework, locate ComputerInit# in the odps-service-computer 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.

- a. 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}
```

- b. Run the following command to check the hostnames in the existing blacklist:

```
/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
```

- c. 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
```

- d. 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.



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.

- a. 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"}
```

Command output:

```
{
  "err_code": 0,
  "err_msg": ""
```

```

    "data": [
      {
        "hostname": "m1"
      }
    ]
  }

```

b. 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 it to the current system time when you configure the audit log. Run the following command to query the mtime:

```

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.

- a. 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 "action_status": "pending" keyword.

- b. 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, and no action indicates that the SR has not been 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 the items in the doing state on the page.

7. Shut down the machine when the status of rma becomes approved or doing. After the maintenance is completed, start the machine.

**Note:**

After the maintenance is completed, if you need to clone the machine, proceed with the next step. Otherwise, skip the next step.

8. Clone the machine.

- a. 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"
```

```
}
]
}
```

- b. Access the clone container. Run the following commands to check the clone status and confirm that 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:1f147fec48
83e082646715cb79c3710f7b2ae9c6e6851fa9a9452b92b4b3366a 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"
```

```
[root@c4 ~]# curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=c43a02002.cloud.a02.amtest1221&attr=state,hostname"
{
  "err_code": 0,
  "err_msg": "",
  "data": {
    "state": "GOOD",
    "hostname": "c43a02002.cloud.a02.amtest1221"
  }
}
[root@c4 ~]#
```

11. Check the cluster final state.

All services on the machine that is brought online have reached the final state.

Server Role	Current Status	Expected Machines	Machines In Final Sta...	Machines Going Offline	Rolling Task Status	Time Used	Actions
checkApsaraAG#	In Final Status	1	1	0	no rolling		Details
checkComputer#	In Final Status	6	6	0	no rolling		Details
checkController#	In Final Status	2	2	0	no rolling		Details
checkCpuStatus#	In Final Status	21	21	0	no rolling		Details
checkPangulMaster#	In Final Status	3	3	0	no rolling		Details
ControllerInit#	In Final Status	1	1	0	no rolling		Details
OdpsController#	In Final Status	2	2	0	no rolling		Details
ServiceTest#	In Final Status	1	1	0	no rolling		Details

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
```

3.2.3.6 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

- In Apsara Infrastructure Management Framework, locate ComputerInit# in the odps-service-computer 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.

- a. 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}
```

- b. Run the following command to check the hostnames in the existing blacklist:

```
/apsara/deploy/rpc_wrapper/rpc.sh blacklist cluster get
```

- c. 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
```

- d. 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

3.2.3.7 Adjust the virtual resources of the Apsara system in MaxCompute

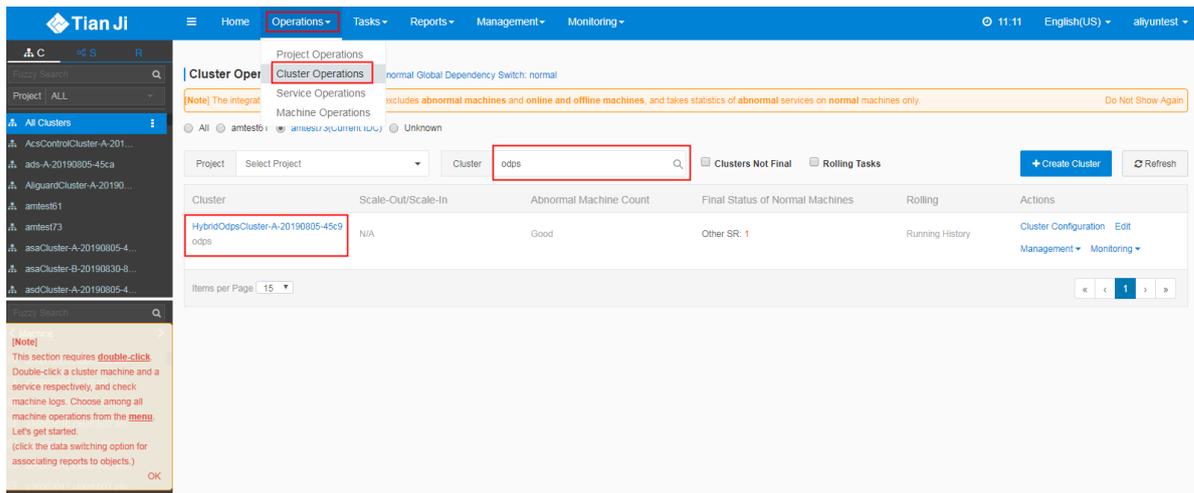
Prerequisites

All MaxCompute services have reached the final state and are functioning properly.

Procedure

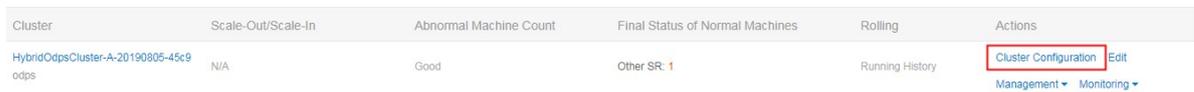
1. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. In the Cluster search box, enter **odps** to search for the expected cluster.

Figure 3-58: Search for a cluster



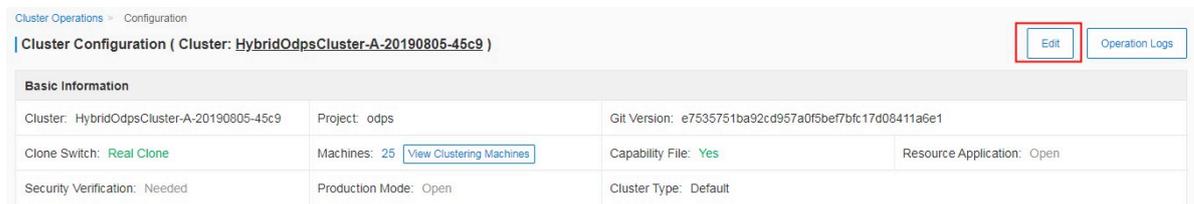
2. Click **Cluster Configuration** to access the cluster configuration page.

Figure 3-59: Cluster Configuration



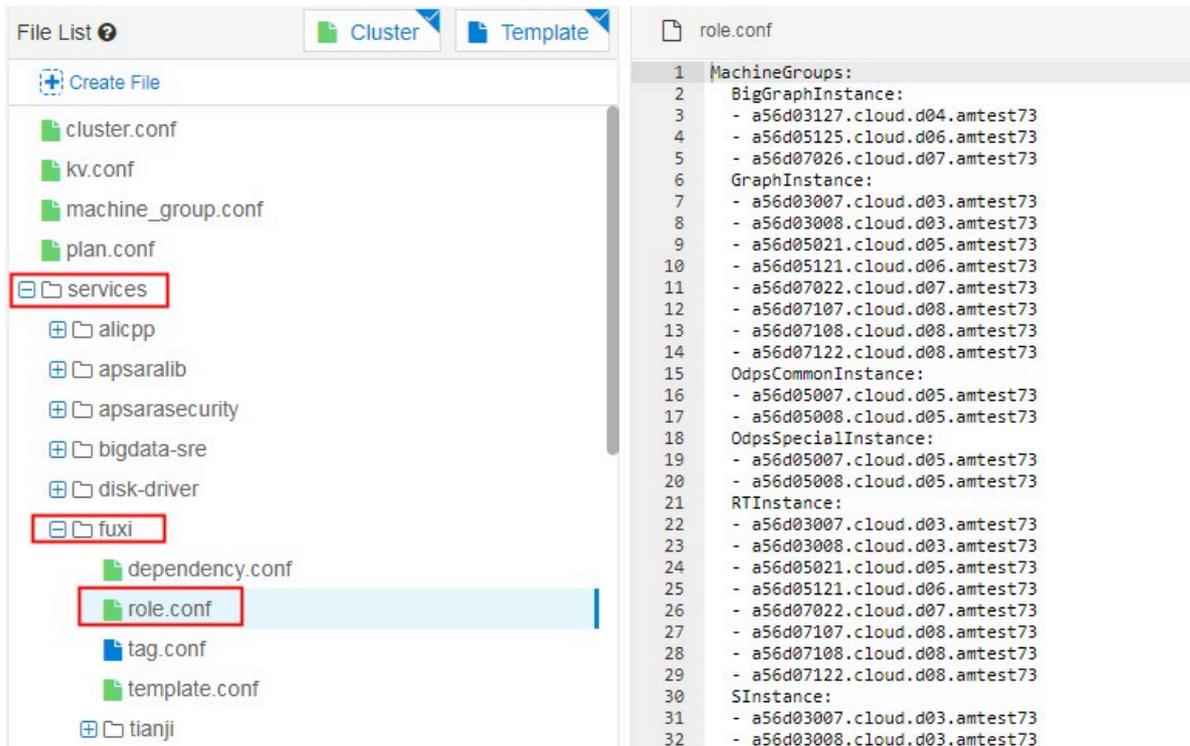
3. On the Cluster Configuration page, click **Edit**.

Figure 3-60: Edit



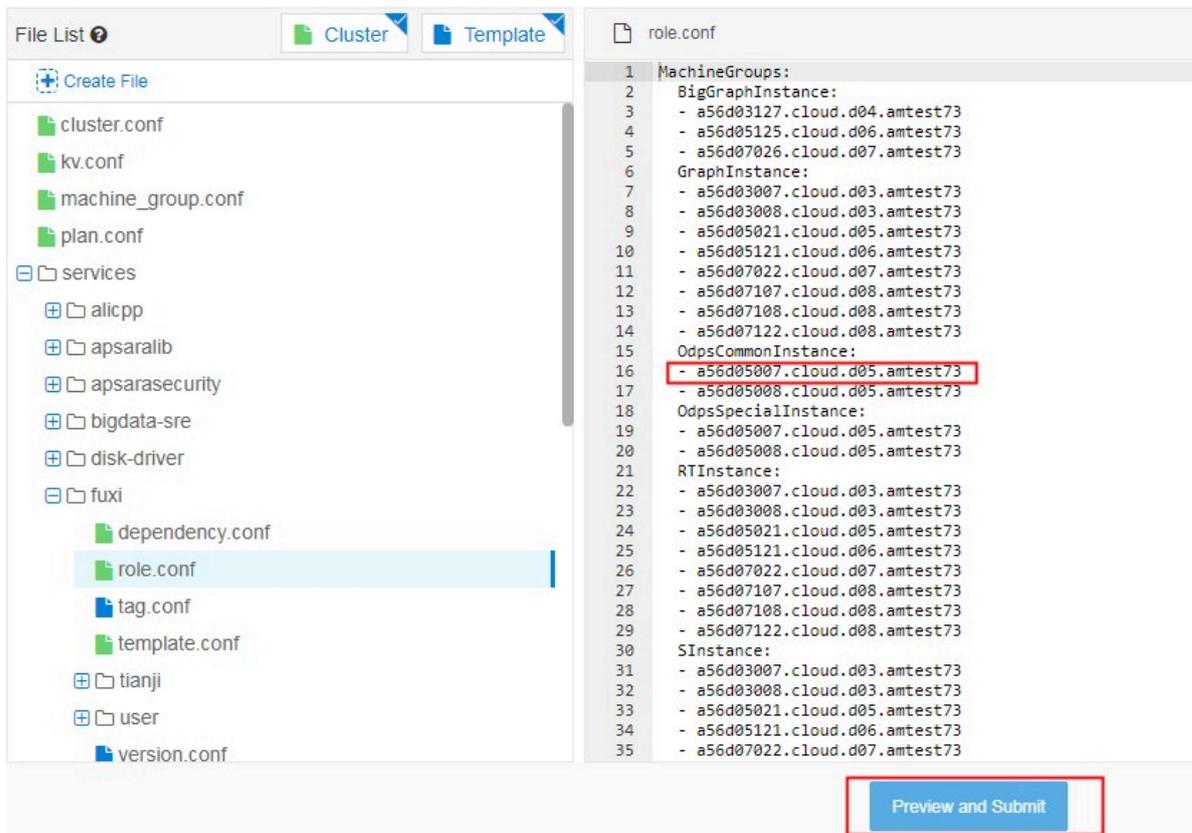
4. In the file list, locate the role.conf file in the fuxi directory.

Figure 3-61: role.conf



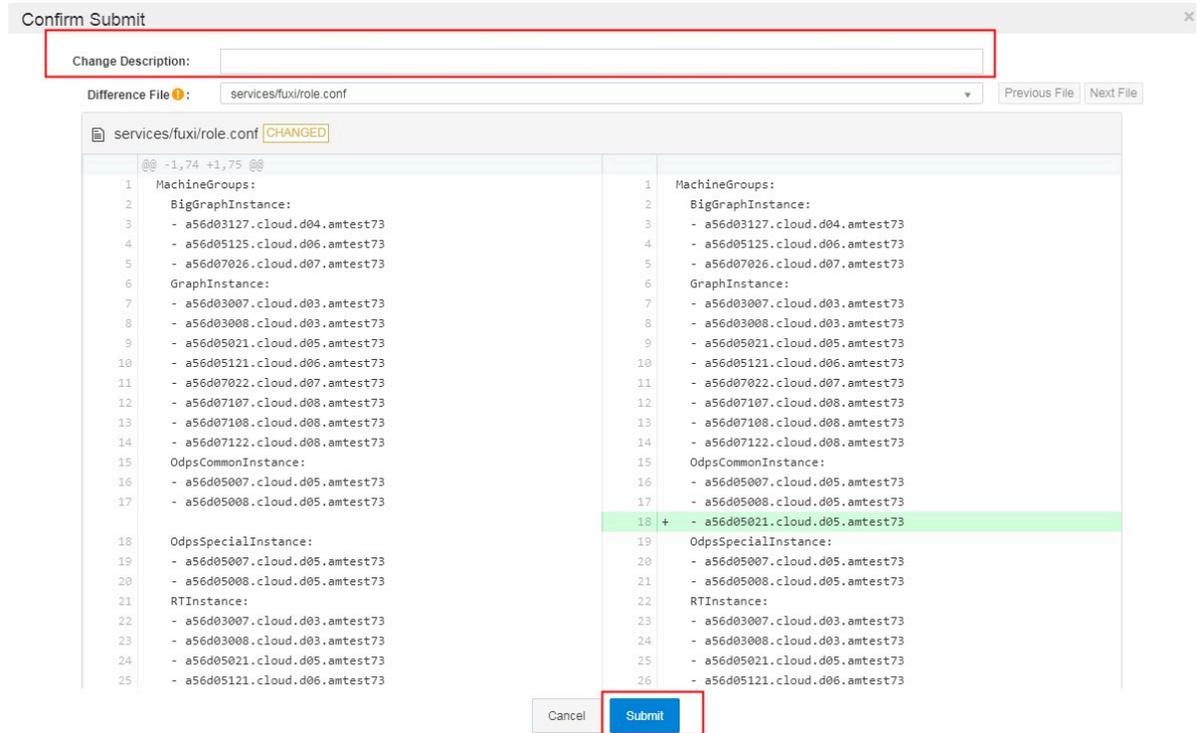
5. Adjust the machine tags on the right and click Preview and Submit.

Figure 3-62: Adjust machine tags



6. In the Confirm Submission dialog box that appears, enter the change description and click Submit.

Figure 3-63: Submit



7. The cluster starts rolling and the changes start to take effect.



Note:

You can check the task status in the operation log. If the status becomes rolling succeeded, it indicates that the changes has taken effect.

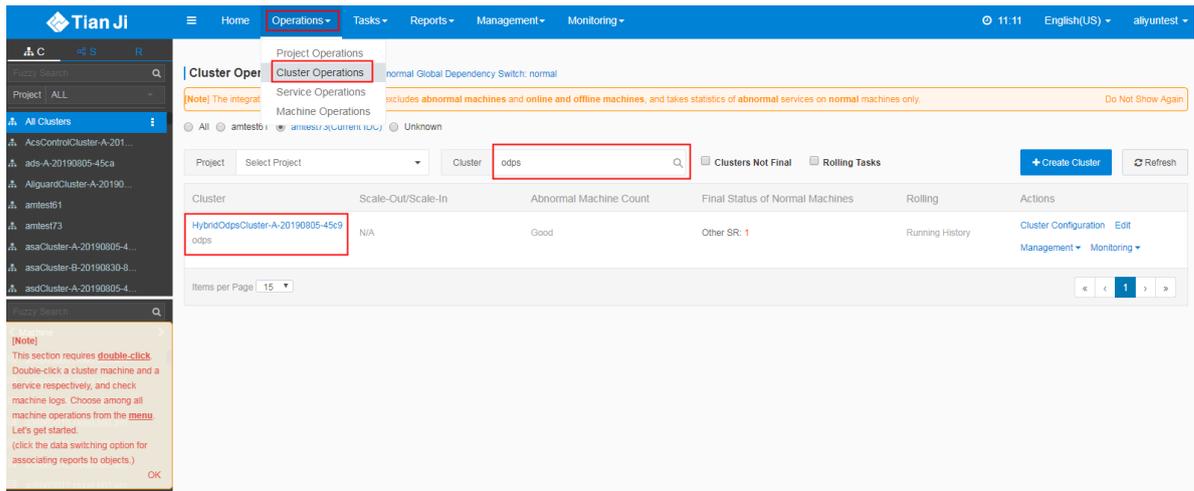
8. After the changes are made, run the `r ttrrl` command in the TerminalService window to confirm the changes.

3.2.3.8 Restart a MaxCompute service

Procedure

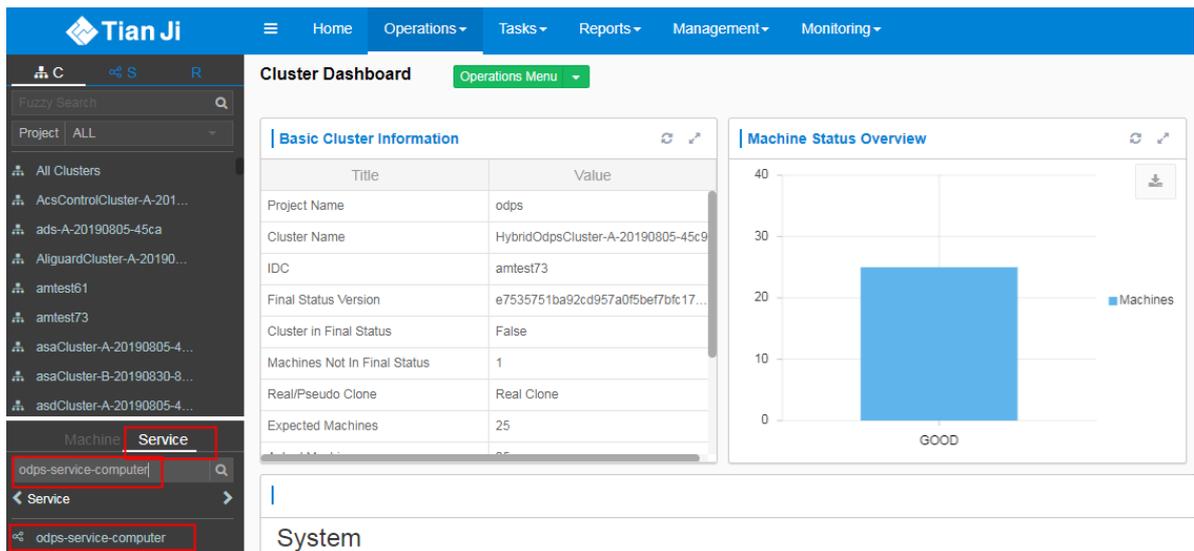
1. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, choose **Operations > Cluster Operations**. In the Cluster search box, enter **odps** to search for the expected cluster.

Figure 3-64: Search for a cluster



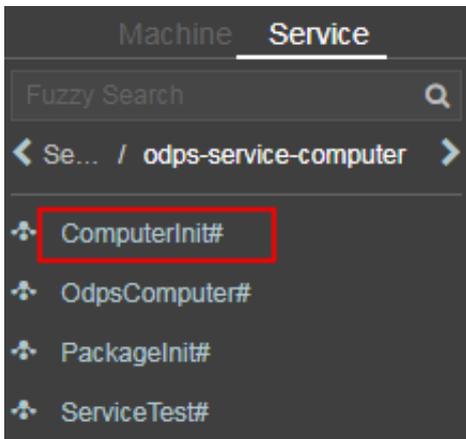
2. Click the cluster in the search result. In the left-side navigation pane, click the **Service** tab, and locate and double-click the **odps-service-computer** service.

Figure 3-65: odps-service-computer



3. After you access the odps-service-computer service, double-click ComputerInit#.

Figure 3-66: ComputerInit#



4. On the page that appears, hover over the vertical dots and choose Terminal from the shortcut menu. In the TerminalService window that appears, you can perform subsequent command line operations.

Figure 3-67: Click Terminal

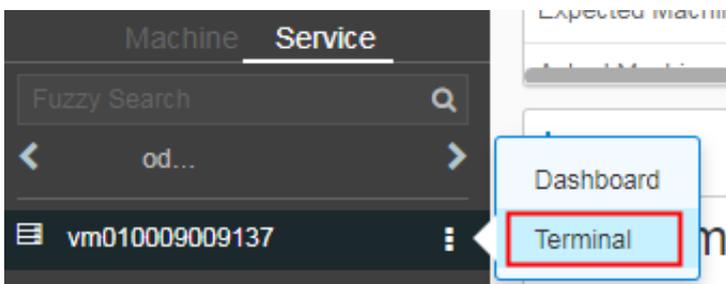
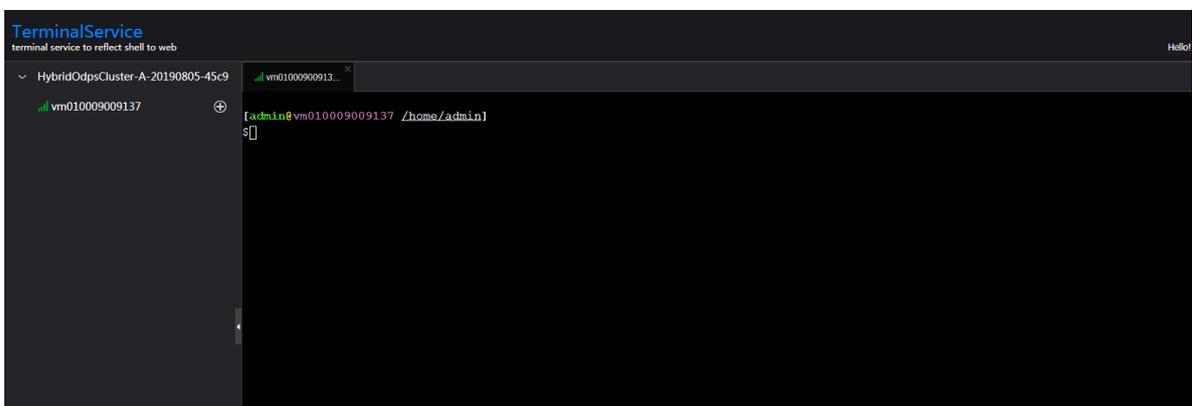


Figure 3-68: Open the TerminalService window



5. Run the following command to obtain the number of machines:

```
tj_show -r fuxi.Tubo#
```

6. Divide the number of machines by 3 to obtain the workernum value.**Note:**

The workersum value ranges from 1 to 3.

7. 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
messenger_partition_num = 2
-- The values here are used as an example. Set these values as
required.
```

8. Restart MaxCompute and Hive.

```
/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.
```

9. Restart the messenger service.

```
cd /apsara/odps_service/deploy/; sh install_odps.sh pedeployme
ssagerservice
-- Restart the messenger service.
```

```
r swl Odps/MessengerServicex
-- Check the service update status and time after restart.
```

10. Restart the quota service.

```
cd /apsara/odps_service/deploy/; sh install_odps.sh pedeployqu
otaservice
-- Restart the quota service.
```

```
r swl Odps/QuotaServicex
-- Check the service update status and time after restart.
```

11. Restart the replication service.

```
cd /apsara/odps_service/deploy/; sh install_odps.sh pedeployre
plicationsservice
-- Restart the replication service.
```

```
r swl Odps/ReplicationServicex
```

```
-- Check the service update status and time after restart.
```

12.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.
```

3.2.4 Common issues and solutions

3.2.4.1 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.

Figure 3-69: Capacity information

```
The pangu 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
```



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.
```

Figure 3-70: Project capacity information

```
$pu ls -l pangu://localcluster/product/aliyun/odps/|grep adsmr -A 4
pangu://localcluster/product/aliyun/odps/adsmr/
Length      : 551267930
FileNumber  : 570
DirNumber   : 143
Pinned      : 0
```



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.

Figure 3-71: 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.

Figure 3-72: 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:

```
pu ls -l pangu://localcluster/product/aliyun/odps/|grep adsmr -A 4
```

Figure 3-73: Number of files for a single project

```
$pu ls -l pangu://localcluster/product/aliyun/odps/|grep adsmr -A 4
pangu://localcluster/product/aliyun/odps/adsmr/
Length      : 551267930
FileNumber  : 570
DirNumber   : 143
Pinned      : 0
```

**Note:****Parameters:**

- **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 ttrtl` command on the cluster AG to view all computing resources.

Figure 3-74: All computing resources

```

$ r ttrtl
total tubo in cluster=13

detail table for every machine:
Machine Name | CPU | Memory | Other
-----
cloud.amtest1284 | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
cloud.amtest1284 | 6,300 | 234,014 | BigGraphInstance:99
cloud.amtest1284 | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
cloud.amtest1284 | 6,300 | 170,453 | ElasticSearchInstance:5
cloud.amtest1284 | 6,300 | 234,014 | BigGraphInstance:99
cloud.amtest1284 | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
cloud.amtest1284 | 6,300 | 170,453 | OdpsSpecialInstance:20 OdpsCommonInstance:20
cloud.amtest1284 | 6,300 | 170,453 | ElasticSearchInstance:5
cloud.amtest1284 | 6,300 | 170,453 | ElasticSearchInstance:5
cloud.amtest1284 | 6,300 | 234,014 | BigGraphInstance:99
cloud.amtest1284 | 6,300 | 170,453 | OdpsSpecialInstance:20 OdpsCommonInstance:20
cloud.amtest1284 | 6,300 | 170,453 | GraphInstance:8 RTInstance:4 SInstance:99
Total | 81,900 | 2,406,572 | 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 Turbo 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.

Figure 3-75: Remaining computing resources

```

$ r tfrl
total tubo in cluster=13
detail table for every machine:
Machine Name | CPU | Memory | Other
-----
cloud. .amtest1284 | 5,025 | 150,990 | GraphInstance:8 RTInstance:4 SInstance:81
cloud. .amtest1284 | 6,090 | 226,874 | BigGraphInstance:98
cloud. .amtest1284 | 5,285 | 153,634 | GraphInstance:8 RTInstance:4 SInstance:83
cloud. .amtest1284 | 6,100 | 68,521 | ElasticSearchInstance:3
cloud. .amtest1284 | 6,190 | 227,850 | BigGraphInstance:98
cloud. .amtest1284 | 6,200 | 169,453 |
cloud. .amtest1284 | 5,035 | 150,450 | GraphInstance:8 RTInstance:4 SInstance:83
cloud. .amtest1284 | 4,600 | 131,565 | OdpsSpecialInstance:15 OdpsCommonInstance:12
cloud. .amtest1284 | 6,200 | 104,921 | ElasticSearchInstance:4
cloud. .amtest1284 | 6,000 | 67,521 | ElasticSearchInstance:3
cloud. .amtest1284 | 5,790 | 218,634 | BigGraphInstance:97
cloud. .amtest1284 | 5,400 | 133,089 | OdpsSpecialInstance:20 OdpsCommonInstance:13
cloud. .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 Turbo machine, as well as the role of each Turbo machine in Job Scheduling System are listed in four columns.

- Run the `r crfu` command on the cluster AG to view the resources used by all running jobs in MaxCompute.

Figure 3-76: Resources used by all running jobs

```

$ r crfu
WorkItemName | CPU | Memory | VirtualResource
-----
odps/DiskDriverService | 280 | 13,600 | {}
odps/odps_elasticsearch_elasticsearch_mdu_es_demo_20170509064623398g2q8q9d | 200 | 1,024 | {}
odps/CGServiceControllerx | 1,980 | 66,660 | {'SInstance': 60}
odps/ReplicationServicex | 200 | 2,000 | {'OdpsSpecialInstance': 1}
odps/OdpsServicex | 1,480 | 45,128 | {'OdpsSpecialInstance': 4, 'OdpsCommonInsta': 7}
odps/HiveServerx | 850 | 37,864 | {'OdpsCommonInstance': 4}
odps/XStreamServicex | 14,070 | 148,370 | {}
odps/QuotaServicex | 160 | 1,024 | {'OdpsSpecialInstance': 1}
odps/MessengerServicex | 300 | 3,092 | {}
sm/sm used resource | 1,900 | 11,192 | {}
total Planned Resource | 20,380 | 327,954 | {'SInstance': 60, 'OdpsSpecialInstance': 11}
    
```



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.

Figure 3-77: 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, ErrorMessage=Storage 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:181)
    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)
Caused 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)
Caused by: RequestId=2018031618244658a751640003a1fa, ErrorCode=InternalServerError, ErrorMessage=Storage quota not enough.
    at com.aliyun.odps.tunnel.io.TunnelRecordWriter.close(TunnelRecordWriter.java:70)
    ... 9 more
ERROR: TunnelException - ErrorCode=Local Error, ErrorMessage=Block ID:0 Failed.
```



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 policies 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.

Figure 3-78: View quota group settings

Account	Alias	SchedulerType	Strategy	InitQuota	ScaledQuota	ScaleRatio	Runtime	UsageInfo
				CPU:31500				CPU:488
			Static		CPU:31500	CPU:37808	CPU:1088	Used
				Mem:852265				Mem:9840
19242	odps_quota	Fair	NoPreempt					
				CPU:168				CPU:488
			Min		Mem:852265	Mem:1022718	Mem:21488	Available
				Mem:1024				Mem:10280

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_cpu_quota $max_mem_quota -m $min_cpu_qu
ota $min_mem_quota
```

 **Note:**

The command with \$QUOTAID is used to modify a quota. The command without \$QUOTAID is used to create a quota.

Figure 3-79: Create a quota

```
$sh /apsara/deploy/rpc_wrapper/rpc.sh setquota -i 9251 -a quotatest -t fair -s 5000 50000 -m 500 500
0
/home/tops/bin/python set_quota_group.py 9251 quotatest 5000 50000 500 5000 fair -1 -1
quotatest
connecting to nuwa://localcluster/sys/fuxi/master/ForClient
connected
Method=SetAccountQuota
Parameter=[{"scaleRatio": {"CPU": 37800, "Memory": 1022718}, "minQuota": {"CPU": 100, "Memory": 1024
}, "returnResourceType": "ReturnResource", "schedulerType": "Fair", "quota": {"CPU": 31500, "Memory"
: 852265}, "canPreemptOtherGroups": false, "canBePreemptedByOtherGroups": false, "alias": "odps quot
a", "strategy": "NoPreempt", "accountId": 9242}, {"scaleRatio": {"CPU": 18900, "Memory": 511359}, "m
inQuota": {"CPU": 100, "Memory": 1024}, "returnResourceType": "ReturnResource", "schedulerType": "Fa
ir", "quota": {"CPU": 18900, "Memory": 511359}, "canPreemptOtherGroups": false, "canBePreemptedByOth
erGroups": false, "alias": "es_quota", "strategy": "NoPreempt", "accountId": 9243}, {"scaleRatio": {
"CPU": 18900, "Memory": 702042}, "minQuota": {"CPU": 100, "Memory": 1024}, "returnResourceType": "Re
turnResource", "schedulerType": "Fair", "quota": {"CPU": 18900, "Memory": 702042}, "canPreemptOtherG
roups": false, "canBePreemptedByOtherGroups": false, "alias": "biggraph_quota", "strategy": "NoPreem
pt", "accountId": 9249}, {"alias": "quotatest", "schedulerType": "Fair", "minQuota": {"CPU": 500, "M
emory": 5000}, "quota": {"CPU": 5000, "Memory": 50000}, "accountId": 9251}]
TraceId=0
TraceLogLevel=ALL
OK
r quota
```

Account/Alias	SchedulerType	Strategy	InitQuota	ScaledQuota	ScaleRatio	Runtime	UsageInfo
			CPU:5000				CPU:0
			Static -----	CPU:5000	CPU:5000	CPU:0	Used -----
			Mem:50000				Mem:0
9251 quotatest	Fair	NoPreempt	-----				
			CPU:500				CPU:0
			Min -----	Mem:50000	Mem:50000	Mem:0	Available -----
			Mem:5000				Mem:0

Figure 3-80: Modify a quota

```
$sh /apsara/deploy/rpc_wrapper/rpc.sh setquota -i 9251 -a quotatest -t fair -s 2000 20000 -m 200 2000
0
/home/tops/bin/python set_quota_group.py 9251 quotatest 2000 20000 200 2000 fair -1 -1
quotatest
connecting to nuwa://localcluster/sys/fuxi/master/ForClient
connected
Method=SetAccountQuota
Parameter=[{"scaleRatio": {"CPU": 5000, "Memory": 50000}, "minQuota": {"CPU": 200, "Memory": 2000}, "returnResourceType": "ReturnResource", "schedulerType":
"Fair", "quota": {"CPU": 2000, "Memory": 20000}, "canPreemptOtherGroups": false, "canBePreemptedByOtherGroups": false, "alias": "quotatest", "strategy": "No
Preempt", "accountId": 9251}, {"scaleRatio": {"CPU": 37800, "Memory": 1022718}, "minQuota": {"CPU": 100, "Memory": 1024}, "returnResourceType": "ReturnResour
ce", "schedulerType": "Fair", "quota": {"CPU": 31500, "Memory": 852265}, "canPreemptOtherGroups": false, "canBePreemptedByOtherGroups": false, "alias": "odps
_quota", "strategy": "NoPreempt", "accountId": 9242}, {"scaleRatio": {"CPU": 18900, "Memory": 511359}, "minQuota": {"CPU": 100, "Memory": 1024}, "returnResou
ceType": "ReturnResource", "schedulerType": "Fair", "quota": {"CPU": 18900, "Memory": 511359}, "canPreemptOtherGroups": false, "canBePreemptedByOtherGroups":
false, "alias": "es_quota", "strategy": "NoPreempt", "accountId": 9243}, {"scaleRatio": {"CPU": 18900, "Memory": 702042}, "minQuota": {"CPU": 100, "Memory":
1024}, "returnResourceType": "ReturnResource", "schedulerType": "Fair", "quota": {"CPU": 18900, "Memory": 702042}, "canPreemptOtherGroups": false, "canBeP
reemptedByOtherGroups": false, "alias": "biggraph_quota", "strategy": "NoPreempt", "accountId": 9249}]
TraceId=0
TraceLogLevel=ALL
OK
r quota
```

Account/Alias	SchedulerType	Strategy	InitQuota	ScaledQuota	ScaleRatio	Runtime	UsageInfo
			CPU:2000				CPU:0
			Static -----	CPU:2000	CPU:5000	CPU:0	Used -----
			Mem:20000				Mem:0
9251 quotatest	Fair	NoPreempt	-----				
			CPU:200				CPU:0
			Min -----	Mem:20000	Mem:50000	Mem:0	Available -----
			Mem:2000				Mem:0

- 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.

Figure 3-81: Jobs waiting for resources

```

admin@docker192168000187 [~/home/admin]
└─$ cruiised -s //g |sort -t | -k2 -rn
sal Planned Resource          9490 243314 ['Instance': 62 'odpsSpecialInstance': 6 'odpsCom
ps.pdata_anc_sit_20180105012108180ghnqgn6_sql_0_1_0_job0 5000 103400 ['Instance': 50]
ps.odpsServicecx             1400 45128   ['odpsSpecialInstance': 4 'odpsCommonInstance': 7]
└─$ sm used resource
ps.HiveServerx               960 35568   ['Instance': 24]
ps.ReplicationServicecx     800 35816   ['odpsCommonInstance': 4]
ps.CGServiceControllerx    400 4000    ['odpsSpecialInstance': 1]
ps.MessengerServicecx     330 11110   ['Instance': 10]
ps.QuotaServicecx           300 3132    ['Instance': 1]
ps.martdata_phq_20180105022552648g822yn_sql_0_1_0_job0 100 1024    ['odpsSpecialInstance': 1]
ps.martdata_lhq_20180105023249573gkbvln6_sql_0_1_0_job0 100 2068   ['Instance': 1]
rkItemName                   CPU   Memory VirtualResource
ps.pdata_lhq_dev_20180105013711904gatuxn_sql_0_1_0_job0 0    0
ps.pdata_lhq_dev_20180105013709696g7wqgn6_sql_0_1_0_job0 0    0
ps.pdata_lhq_dev_20180105013544194gwsuxn_sql_0_1_0_job0 0    0
ps.pdata_lhq_dev_20180105013401776ggsuxn_sql_0_1_0_job0 0    0
ps.pdata_lhq_dev_20180105013237183g7rufln6_sql_0_1_0_job0 0    0
ps.pdata_anc_dev_201801050223353629812yn_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_20180105022546325gdefen6_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_20180105022545896g122yn_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_20180105022545533g022yn_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_20180105022544217gcefen6_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_20180105022537950g79vln6_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_20180105022535810gvl2yn_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_2018010502253566g59vln6_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_20180105022535487g49vln6_sql_0_1_0_job0 0    0
ps.odata_lhq_sit_20180105022534509gt12yn_sql_0_1_0_job0 0    0
ps.odata_lhq_dev_20180105022351607g68vln6_sql_0_1_0_job0 0    0
ps.odata_lhq_dev_20180105013544819gvpqgn6_sql_0_1_0_job0 0    0
ps.odata_ghq_sit_20180105014404120g7vuxn_sql_0_1_0_job0 0    0
ps.odata_ghq_sit_20180105011745244gskuxn_sql_0_1_0_job0 0    0
ps.martdata_phq_dev_20180105021429726glafen6_sql_0_1_0_job0 0    0
ps.martdata_phq_dev_20180105014324311g9yqgn6_sql_0_1_0_job0 0    0
ps.martdata_lhq_dev_2018010502110917gslrgn6_sql_0_1_0_job0 0    0
ps.martdata_lhq_dev_20180105014316771g8yqgn6_sql_0_1_0_job0 0    0
admin@docker192168000187 [~/home/admin]
    
```

Figure 3-82: Quota used up

```

admin@docker192168000187 [~/home/admin]
└─$ quota
-----
account|Alias          |SchedulerType |Strategy |InitQuota          |ScaledQuota      |ScaledRatio      |Runtime      |UsageInfo
-----|-----|-----|-----|-----|-----|-----|-----|-----
9242  |odps_quota        |Fair          |NoPreempt|Static             |CPU:42000        |CPU:42000        |CPU:0        |Used      |CPU:0
                                                |Mem:1293336     |Mem:1293336     |Mem:0        |Mem:0
                                                |CPU:100         |Mem:343489      |Mem:0        |Mem:0
                                                |Mem:1024
9243  |kaifa             |Fair          |NoPreempt|Static             |CPU:5000         |CPU:5000         |CPU:5000    |Used      |CPU:5000
                                                |Mem:620886     |Mem:620886     |Mem:0        |Mem:103400
                                                |CPU:100         |Mem:164506     |Mem:0        |Mem:0
                                                |Mem:100
9244  |phq               |Fair          |NoPreempt|Static             |CPU:42000        |CPU:12370        |CPU:42000   |CPU:100   |Used      |CPU:0
                                                |Mem:1293336     |Mem:342565      |Mem:1293336|Mem:2068
                                                |CPU:100         |Mem:100         |Mem:0        |Mem:0
9245  |lhq               |Fair          |NoPreempt|Static             |CPU:42000        |CPU:12370        |CPU:42000   |CPU:0     |Used      |CPU:0
                                                |Mem:1293336     |Mem:342565      |Mem:1293336|Mem:0
                                                |CPU:100         |Mem:100         |Mem:0        |Mem: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.

3.2.4.2 Common issues and data skew troubleshooting

Scenario 1: how to determine whether a job stops running because of insufficient resources

Symptom: The job does not progress as expected.

Figure 3-83: 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 indicates that there are insufficient resources to run these tasks. After an instance obtains the necessary resources, its status changes to Running.

	FuxiInstanceID	IP & Path	StdOut	StdErr	Status
1	Odps/odps_s...				Ready
2	Odps/odps_s...				Ready
3	Odps/odps_s...				Ready
4	Odps/odps_s...				Ready
5	Odps/odps_s...				Ready

Solution:

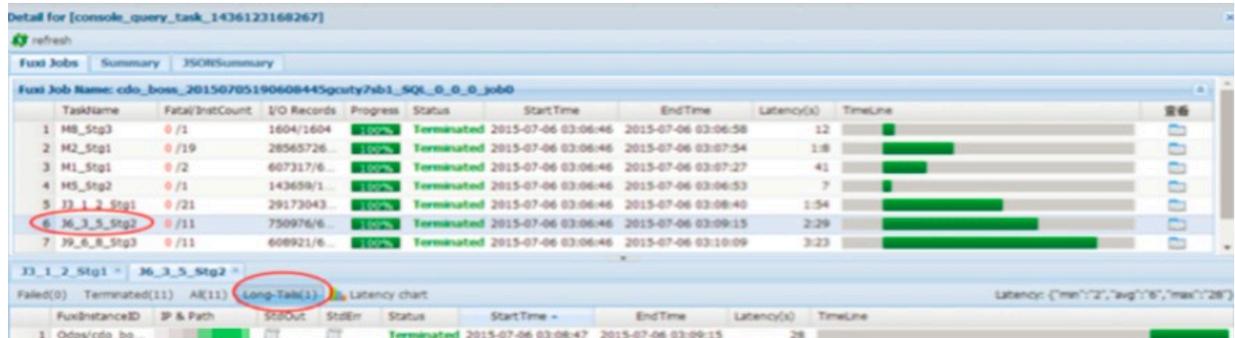
- If there are insufficient resources during peak hours, you can run the tasks during off-peak hours.
- If there are insufficient computing quotas, check whether there are sufficient computing resources in the quota group of the project.
- If computing resources in the cluster are occupied for long periods of time, you can develop a computing quota allocation policy to scale the quota.
- 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 FIFO scheduling policy.

Scenario 2: how to find the root cause for the extended running period of a job

Symptom: The MaxCompute job execution progress remains 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.

Figure 3-84: Cause analysis



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.

Figure 3-85: Further analysis

```

R2_1_Stg1:
  instance count: 1
  run time: 12.000
  instance time:
    min: 0.000, max: 0.000, avg: 0.000
  input records:
    input: 15 (min: 15, max: 15, avg: 15)
  output records:
    R2_1_Stg1FS_11934: 15 (min: 15, max: 15, avg: 15)

```

Solution: If slow Fuxi instances exist on a particular machine, analyze whether the machine has a hardware failure.

Scenario 3: how to improve the concurrency of MaxCompute jobs

Fault location: 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 one partition based on the `odps.sql.mapper.merge.limit.size` value and be processed by one instance. This

improves instance utilization. The `odps.sql.mapper.merge.limit.size` value is 64 MB by default, and the total size of small files merged will not exceed this value.

- Instances cannot process data across multiple partitions.

A partition is mapped to a folder in Apsara Distributed File System. You need to run at least one instance to process the 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 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 based on the split size of 256 MB, a single Map instance needs to process a large number of records. This reduces the record processing concurrency.

- Dump occurs in the Map, Reduce, and Join stages.

Based on the preceding analysis of job summary, the display of dump information indicates that the instance does not have sufficient memory for data sorting in the Shuffle stage. Improving the concurrency can reduce the amount of data processed by a single instance to the amount of data that can be handled by the memory. This eliminates disk I/O time consumption and improves the processing speed.

- Time-consuming UDFs are used.

The execution of UDFs is time-consuming. If you execute UDFs concurrently, you can reduce the processing 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 consume more resources. We recommend that you take cost into account when improving concurrency. After optimization, an instance takes an average of 10 minutes to complete, 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, run 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, filter out the null values and use a case when statement to add a random number to each key before the Join operation.
- If you do not want to change 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/false
```

- **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 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 to process each Map task instance. **Default value:** 100. **Valid values:** 50 to 800.

```
set odps.sql.mapper.memory=1024
```

Description: allows you to set the memory size of each Map task instance. **Unit:** MB. **Default value:** 1024. **Valid values:** 256 to 12,288.

```
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 Join task instances. **Default value:** -1. **Valid values:** 0 to 2,000.

```
set odps.sql.joiner.cpu=100
```

Description: allows you to set the number of CPUs used to process each Join task instance. **Default value:** 100. **Valid values:** 50 to 800.

```
set odps.sql.joiner.memory=1024
```

Description: allows you to set the memory size of each Join task instance. **Unit:** MB. **Default value:** 1024. **Valid values:** 256 to 12,288.

Reduce settings

```
set odps.sql.reducer.instances=-1
```

Description: allows you to set the number of Reduce task instances. **Default value:** -1. **Valid values:** 0 to 2,000.

```
set odps.sql.reducer.cpu=100
```

Description: allows you to set the number of CPUs used to process each Reduce task instance. **Default value:** 100. **Valid values:** 50 to 800.

```
set odps.sql.reducer.memory=1024
```

Description: allows you to set the memory size of each Reduce task instance. **Unit:** MB. **Default value:** 1024. **Valid values:** 256 to 12,288.

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 12,288.

```
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 3,600.

```
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 3,072.

```
set odps.sql.udf.optimize.reuse=true/false
```

Description: When this parameter is set to true, each UDF function expression can be calculated only 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 2,048.

```
set odps.sql.reshuffle.dynamiccpt=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 avoid 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.

Figure 3-86: Storage information

```

odps@ odps_smoke_test>desc project odps_smoke_test -extended;
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
tableLogicalSize                      20530
usedQuotaPhysicalSize                 4162347
resourcePhysicalSize                  4043403
tempResourcePhysicalSize              0
tableBackupPhysicalSize               38016
volumePhysicalSize                   0
volumeLogicalSize                    0
failoverPhysicalSize                  8412
tableBackupLogicalSize                12672
failoverLogicalSize                   2804
tempResourceLogicalSize              0
tablePhysicalSize                     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.

3.3 DataWorks

3.3.1 Basic concepts and structure

3.3.1.1 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).

3.3.1.2 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.

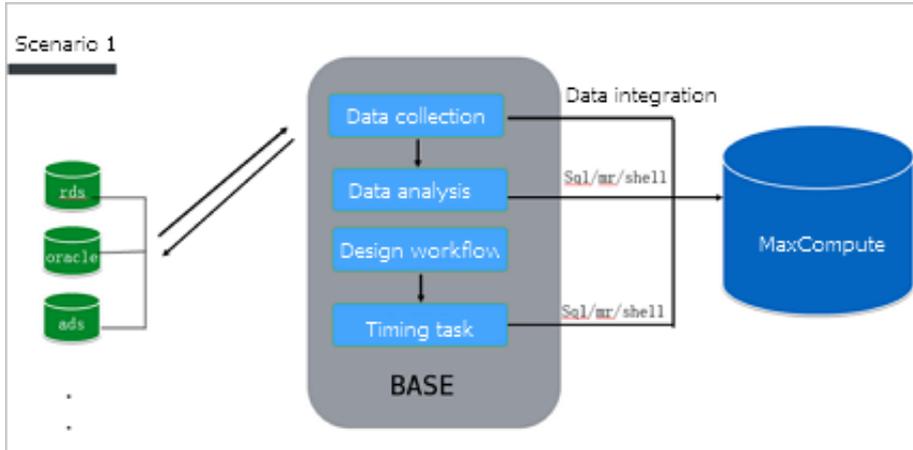
3.3.1.3 Introduction to data analytics

Scenario 1: data synchronization and analysis

Scenario 1 shows an example of data analytics in typical scenarios.

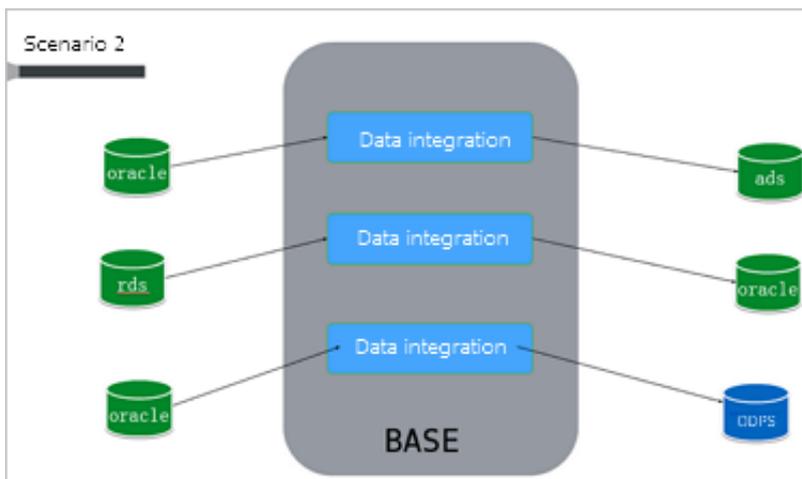
- 1. You can collect data from various databases, and send the data to MaxCompute by using DataWorks.**
- 2. You can log on to DataWorks, create SQL, MapReduce, and shell nodes, and commit the nodes to MaxCompute for data analysis.**

3. You can use DataWorks to synchronize the analysis results from MaxCompute to the databases from which you collect data.



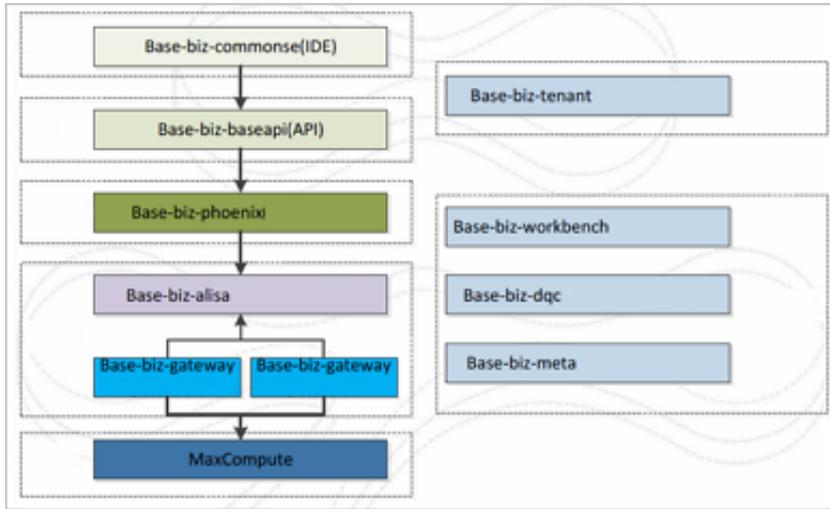
Scenario 2: data synchronization

DataWorks supports data synchronization between various databases. You can synchronize data by using DataWorks.



3.3.1.4 Architecture of DataWorks in Apsara Stack V3

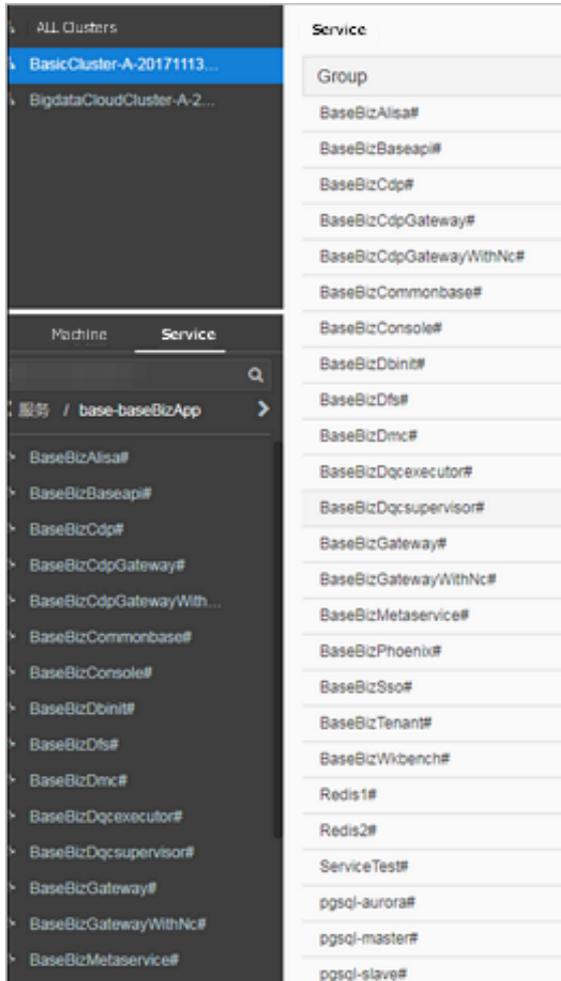
Figure 3-87: The core architecture of base



Services in DataWorks play an important role for node scheduling and running. You can perform all O&M operations for DataWorks of Apsara Stack V3 in the Apsara

Infrastructure Management Framework. DataWorks consists of the following services.

Figure 3-88: base components



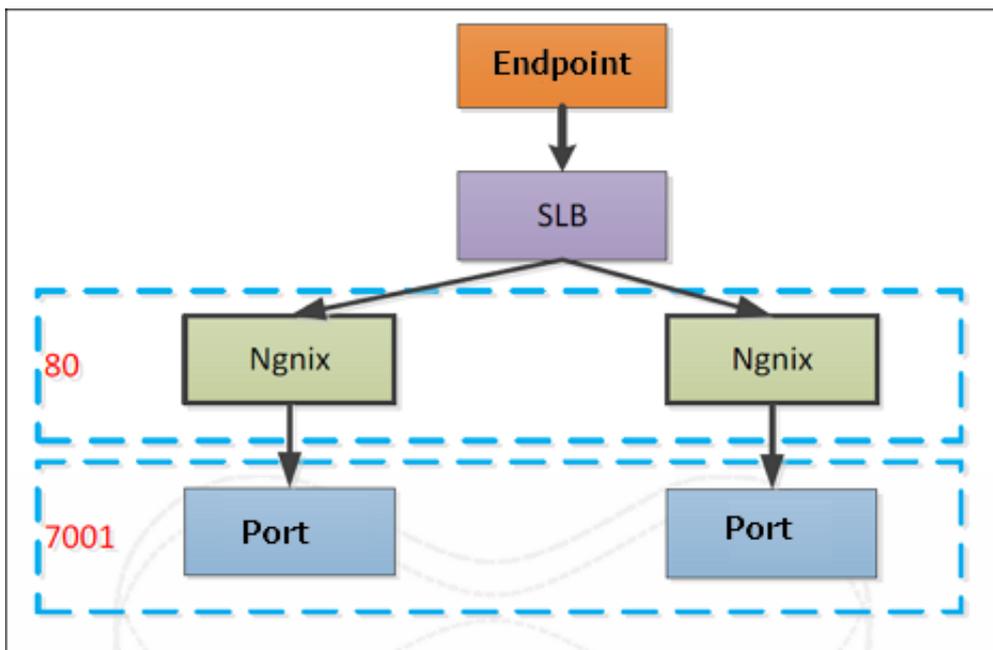
All services in DataWorks are deployed on Docker containers. You can log on to a host, and run the docker ps command to view the containers on which the services are deployed.

```

[admin@base ~]$ docker ps
CONTAINER ID        IMAGE               COMMAND             CREATED             STATUS              PORTS
3d7704678d8        "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8014->80/tcp
c1e1809089         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8013->80/tcp
8d758273d9         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8012->80/tcp
38646d2775f        "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8011->80/tcp
f0c488494e         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8010->80/tcp
31c125456b14         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8009->80/tcp
6d8094623c3         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8008->80/tcp
8c71a4741332         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8007->80/tcp
897140794470         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8006->80/tcp
...
113126102ff4         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8005->80/tcp
8b717608112b         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 80/tcp, 0.0.0.0:8004->7001/tcp
319a287758fa         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8003->80/tcp
11793829254         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8002->80/tcp
31c48142d4d5         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:8001->80/tcp
3405ab255b1         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:80->80/tcp
2ae57c94f48         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 6379/tcp, 0.0.0.0:36379->36379/tcp
688487cc935         "/bin/bash -c /start..." "/bin/bash -c /start..." 3 months ago        Up 3 months        22/tcp, 0.0.0.0:16379->16379/tcp, 6379/
48315ab5886         "/bin/bash -c /start..." "/docker-entrypoint.s..." 3 months ago        Up 3 months        0.0.0.0:5432->5432/tcp
4833915c2187         "/bin/bash -c /start..." "/entrypoint.sh mysq..." 3 months ago        Up 3 months        0.0.0.0:3306->3306/tcp
674888c83354         "/bin/bash -c /start..." "/bin/sh -c 'usr/loa..." 3 months ago        Up 3 months        80/tcp, 0.0.0.0:1022->12/tcp
76d78e4a21cd         sware              "/swarm json --advert..." 4 months ago        Up 4 months        2375/tcp
    
```

The internal structure of each service (except gateway) is shown in *Figure 3-89: Internal structure*.

Figure 3-89: Internal structure



3.3.1.5 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.

3.3.2 Common administration tools and commands

3.3.2.1 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.

3.3.2.2 Cluster resource list

In the Apsara Infrastructure Management Framework, select **base** from the project drop-down list, select **BasicCluster**, move the pointer over the **More** icon next to **BasicCluster**, and select the **Dashboard** option to open the Cluster Dashboard of the **BasicCluster**.

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.

3.3.2.3 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/alisa/tasknode/target/alisa/tasknode/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 restart`.

3.3.2.4 View the log of a failed task

If you see a failed task in Administration, click **Failed** to view all failed tasks.

Find a task, and choose **View Log** from the **More** drop-down list.

3.3.2.5 Rerun a task

If you want to rerun a failed task, select the task in the Administration console and click **Rerun**.

3.3.2.6 Terminate a task

If you want to terminate a running task, select the task in Administration, and then click **Terminate**.



Note:

Only running tasks can be terminated.

3.3.2.7 Filter tasks in the administration center

You can choose **Administration > Task List** and filter the tasks to maintain.

3.3.2.8 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]
```



Note:

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 view the container logs.

`docker exec -it [container ID] bash`: You can run this command to enter the container.

3.3.2.9 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;
```

3.3.3 Process daily administration operations

3.3.3.1 Daily check

3.3.3.1.1 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.

3.3.3.1.2 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' order by time desc;
```

In the result, if the stat value is active, the service is running properly. If not, contact Alibaba Cloud Customer Support.

3.3.3.1.3 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.

3.3.3.1.4 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.

3.3.3.2 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`.

3.3.3.3 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****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 Infrastructure 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.conf 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.conf 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.conf 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.conf 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 [machine1,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 [machine1,machine2] --config clt2.conf
```

Parameters

- **-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 `tail -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 `select * from 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.

3.3.3.4 Scale in the node cluster that runs the base-biz-gateway service

Prerequisites

If a server that runs the base-biz-gateway service fails, you can repair and restart the server to redeploy the server.

If the server that runs the base-biz-gateway service functions as expected, you can take the server out of service and then move it from the BasicCluster cluster to the default cluster by following the instructions provided in this document.



Note:

Risks:

- Before taking the server out of the base-biz-gateway service, make sure that no applications are running on the server.
- You also need to make sure that the hostname of the server is correct.

Procedure

Prepare for the cluster scale-in

1. Collect server and cluster information.

You need to collect the detailed information of the target server and the cluster to which the target server belongs.

2. You need to make sure that the value of the scalable tag is true for the service resource group (SRG) to which the server belongs. If the value is false, you need to change it to true and submit a rolling task.

On the left-side navigation pane, choose BasicCluster > Cluster Configuration File > machine_group.conf. In this file, you need to verify that the value of the scalable tag is true for the server to be removed from the cluster.

Stop the gateway service

1. Log on to the target server, and run the `ps -ef|grep gateway` command to obtain the container ID.
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/serverctl stop` command.
4. Run the `ps -ef|grep java` command to check whether any process is running on the server. If a process is running, run the `kill -9 [pid No.]` command to terminate the process.
5. Delete the program directories from the server.

Clean up the disks of the server. You can skip this step if you need to clone the server.

```
#rm -rf /home/admin/*
```

```
#rm -rf /opt/taobao/tbdpapp/
```

Move a server from a node cluster to the default cluster in the Apsara Infrastructure Management Framework

1. Run the following commands on the ops1 server to remove a server:

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf (You can ignore the
message indicating that the symbolic link already exists.)
./tianji_ops_tool.py contract_nc -c [clusterName] -l [machineList]
--config tianji_clt.conf -s [SRGname]
```

Parameters

- **-c:** the name of the BasicCluster cluster. Select base from the project drop-down list, and then you can find the cluster name. This parameter is required.
 - **-l:** the hostnames of servers to be removed from the cluster. Separate multiple hostnames by commas (.). This parameter is required.
 - **-s:** the name of the SRG where the servers reside. This parameter is required. You can open machine_group.conf in the BasicCluster configuration file list. In the file, you can find the SRG where the servers reside.
 - **-config:** the tianji_clt configuration file. This parameter is required.
2. After you run the command, check whether the scale-in operation succeeds in the Apsara Infrastructure Management Framework.

Go to the Cluster Operation and Maintenance Center of the BasicCluster cluster.

3. On the Cluster Operation and Maintenance Center page, you can click the number next to Machine: in: to identify the status of the servers that are to be removed.

If the scale-in operation succeeds, the number of servers that are being taken out of service decreases to zero. Otherwise, you need to 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 the Apsara Infrastructure Management Framework. If you need to remove the servers from the Apsara Infrastructure Management Framework, proceed as follows.

Remove a server from the Apsara Infrastructure Management Framework

1. Choose Operations > Machine Operations.
2. On the Machine Operations page, select Machine Online/Offline in the upper-right corner.
3. Select the server that you want to remove from the default cluster.

4. Search for the server that you want to remove by hostname on the left-side navigation pane. You can only remove servers in the default cluster.
5. If the status of the server is Offline Ready, click Confirm Offline.

Verify the server removal result

1. You can verify that the server is moved to the default cluster in the Apsara Infrastructure Management Framework.

Choose Operations > Machine Operations. Search for the server by its hostname, and check whether it is in the default cluster.

2. Verify that the server is removed from the default cluster.

Choose Operations > Machine Operations page, and search for the server by its hostname. If you cannot find the server in the search results, the server is removed.

3. Run the following command on the ops1 server to check whether the server is removed from the default cluster by calling an API operation.

```
curl http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=$hostname
```

3.3.3.5 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.
3. 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 alisa` 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.

3.3.3.6 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.
2. 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 phoenix` 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.

3.3.3.7 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.

3.3.3.8 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 base-baseBizApp, 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/bin/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.

3.3.3.9 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/bin/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.

3.3.3.10 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 base-baseBizApp, 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
```

3.3.3.11 Restart DataWorks Data Service

Procedure

1. In the Apsara Infrastructure Management Framework console, search for **dataworks-dataservice** 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.

3.3.3.12 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.

3.3.4 Common issues and solutions

3.3.4.1 Node instances remain in the Pending (Resources) status

Description

Excessive tasks to be run on the BasicCluster cluster remain in waiting status. In Operation Center, you may find that many instances of recurring nodes remain in the Pending (Resources) status for a long period of time.

Cause

The issue may occur because of one of the following four reasons:

- The gateway server is overloaded or offline.
- The service reaches the maximum concurrency.
- The gateway server runs out of disk space.
- The NTP time of servers in the BasicCluster cluster is inconsistent.

Solution

To resolve this issue, perform the following steps:

- Check the status of the gateway server in the database.
 1. Log on to the database that hosts the base-biz-alisa service. In Apsara Stack V3 , you can find the database address from the resource list of the BasicCluster cluster in the Apsara Infrastructure Management Framework.
 2. Run the `select * from alisa_node;` command to check the values in the `active_type` and `live` fields.

If the live field value is -1, the server is offline. If the active_type field value is -1, the server is overloaded.



Note:

In either case, you need to 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 function run is running expected. Otherwise, you need to check the configuration file.

- Run the `top` command to display the load.

The status of the server is -1 as the result of high load. In this case, you need to check the top command output, and see whether any node causes high CPU and memory usage.

You can run the `ps -ef|greppid` command to view processes of the specified node, and identify which process causes the high load. Run the `kill -9 pid` command to terminate the process of the node. After the load drops, check whether the status of the server changes to 1.

- Check whether the base-biz-gateway service reaches the maximum concurrency.

Log on to the alisa database, and run the following SQL 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;
```

You can compare the query results of the two statements.

- The query result of the first statement displays the maximum concurrency of each resource group.
- The query result of the second statement displays the current concurrency status of each resource group.

If the query results of the two statements are the same or similar, the service reaches the maximum concurrency. You need to make sure that available resources are sufficient before running a node.

You can run the following SQL statement to list the 10 nodes that require the longest runtime.

```
select task_id,gateway,slot,create_time from alisa_task where status
=2 and create_time>current_time order by create_time desc limit 10;
```

You can 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 task IDs that are returned by the preceding SQL statement. You can obtain the task name from the command output.

Then, you can troubleshoot the node. You can run the `kill -9` command to terminate the node and release resources if required. Otherwise, if you want to run new nodes, you have to wait until these existing nodes have finished.

- Check whether the disk on the 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, you can run the `du -sh` command and identify the files under the `/home/admin` directory that consume a large amount of space. If large log files exist in the `/home/admin/alisatasknode/taskinfo/` directory, you can remove them.

- Check the NTP time of servers in the BasicCluster cluster.
 1. Log on to the alisa database and run the `select now();` command to view the current time of the database.
 2. Check the system time of hosts in the BasicCluster clusters against the time in the alisa database.
 3. Run the `date` command on the hosts to check whether the system time of each host is consistent with the time in the alisa database. If the discrepancy is greater than 30s, the service may fail and you need to adjust the NTP time of the hosts in the BasicCluster cluster.

**Note:**

In Apsara Stack V3, you need to locate the servers in the service list and follow the proceeding steps to resolve the issue.

- Rename the directory, and then restart the service.

If the issue persists after you perform these four steps, you can 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. The format is `YYYYMMDD`.

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 under the directory has reached the maximum and you cannot create more subdirectories. To resolve this issue, perform the following steps:

1. Rename the `/home/admin/alisatasknode/taskinfo/20180306/phoenix` directory as `/home/admin/alisatasknode/taskinfo/20180306/phoenix.bak`.
2. Run the following command to restart the service:

```
sudo su admin -c "/home/admin/alisatasknode/target/alisatasknode/bin/serverctlrestart"
```

**Note:**

This issue may occur when the gateway server uses the `ext3` file system.

3.3.4.2 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='reader';
```

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='reader';
```

3.3.4.3 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,"opSEQ":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.

3.3.4.4 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.

3.3.4.5 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.

3.3.4.6 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 {} \;
```

3.3.4.7 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.

3.4 Realtime Compute

3.4.1 Job status

3.4.1.1 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.

3.4.1.2 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.

3.4.1.3 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.

3.4.1.4 Job instantaneous values

Table 3-10: Job parameters

Name	Description
Consumed compute time	Indicates the computing performance of a job.
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.

Name	Description
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.

3.4.1.5 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:

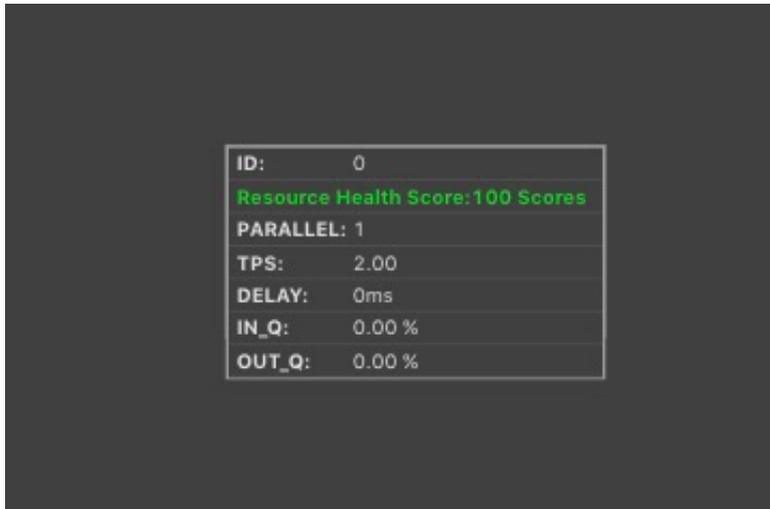
- Reduces the thread-to-thread handover.
- Reduces the message serialization and deserialization.
- Reduces the data handover in the buffer zone.
- Increases overall throughput while decreasing latency.

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 [Figure 3-90: View mode](#), to offer you a more intuitive display.

Figure 3-90: View mode



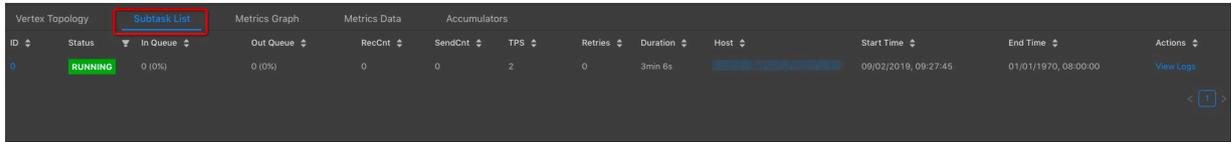
You can view the detailed information about a task by moving the pointer over the task. [Table 3-11: Parameter description](#) describes the task parameters.

Table 3-11: Parameter description

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.
MEM	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.

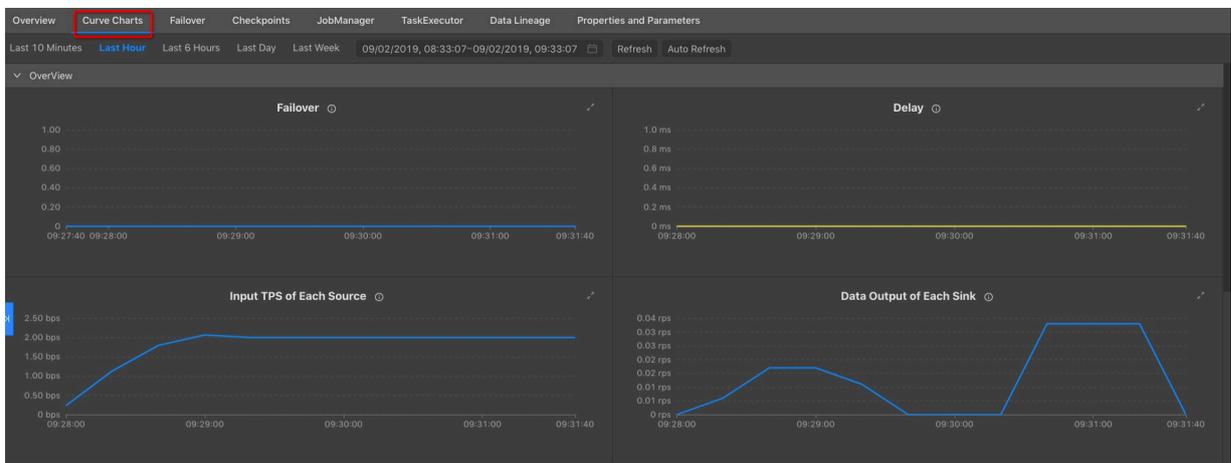
You can also click a task node to access its details page. On this page, you can view its subtasks, as shown in [Figure 3-91: Task details page](#).

Figure 3-91: Task details page



The Curve Charts tab provides curve charts to show the metrics of each task, as shown in [Figure 3-92: Curve charts for task metrics](#).

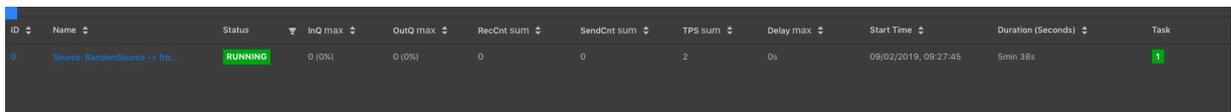
Figure 3-92: Curve charts for task metrics



List mode

In addition to the view mode, Realtime Compute also allows you to view each task in the list mode, as shown in [Figure 3-93: List mode](#).

Figure 3-93: List mode



[Table 3-12: Parameter description](#) describes the task parameters.

Table 3-12: Parameter description

Parameter	Description
ID	The task ID in the running topology.

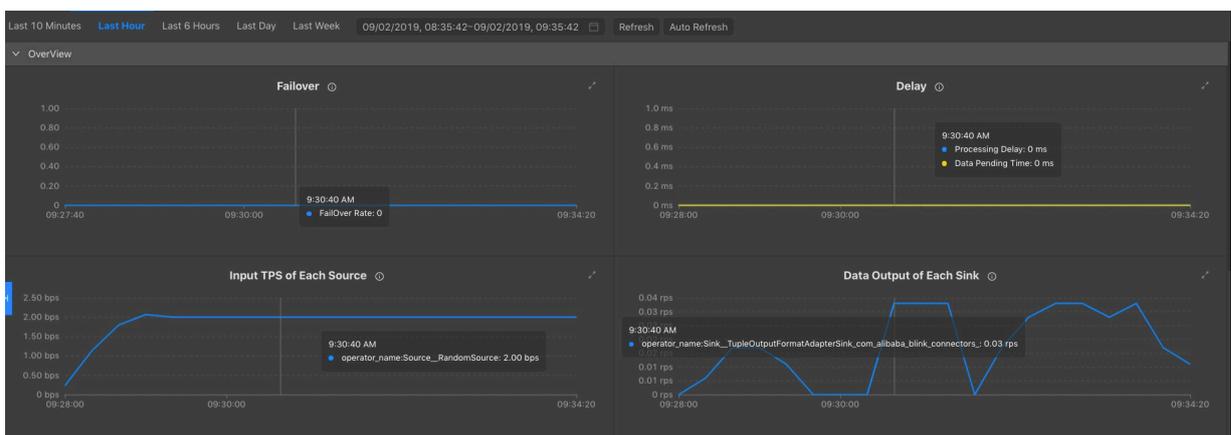
Parameter	Description
Name	The name of the task.
Status	The status of the task.
INQ max	The maximum percentage of input queues for the task node.
OUTQ max	The maximum percentage of output queues for the task node.
RecvCnt sum	The total amount of data that is received by the task node.
SendCnt sum	The 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.

3.4.2 Curve charts

3.4.2.1 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.

Figure 3-94: Curve Charts tab



**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.

3.4.2.2 Overview

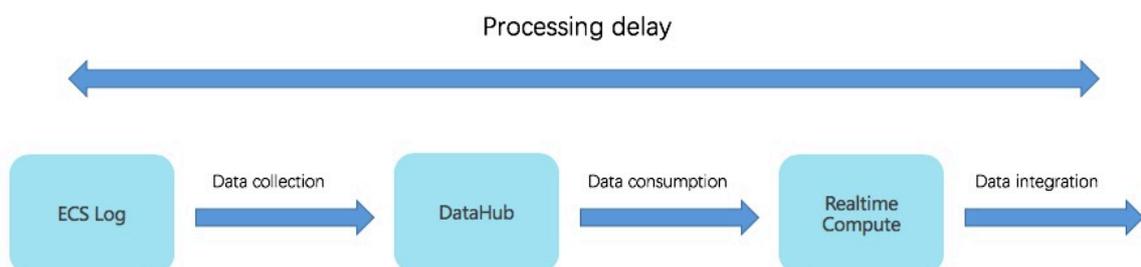
Failover rate

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 increases accordingly. The following table shows the processing delay.

Figure 3-95: 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.



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

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.

3.4.2.3 Advanced view

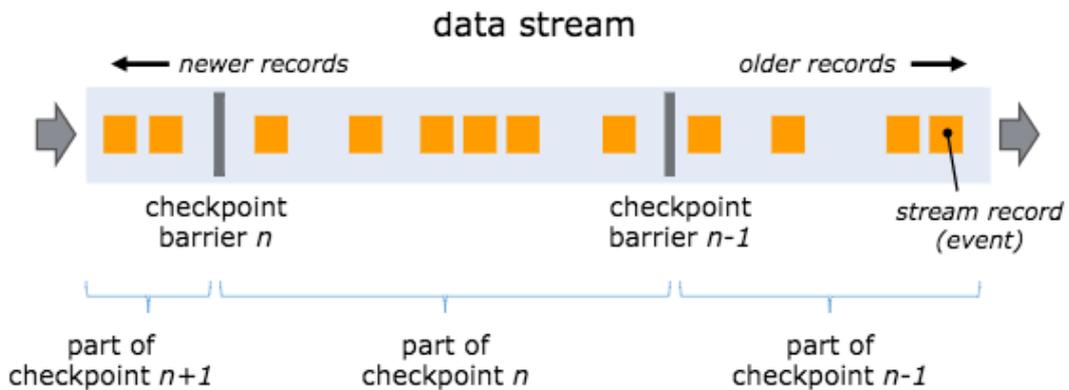
Realtime Compute offers a fault tolerance mechanism to consistently recover the state of data streaming applications. The central part of the fault tolerance mechanism is drawing consistent snapshots of the distributed data stream and the state. These snapshots act as consistent checkpoints to which 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.

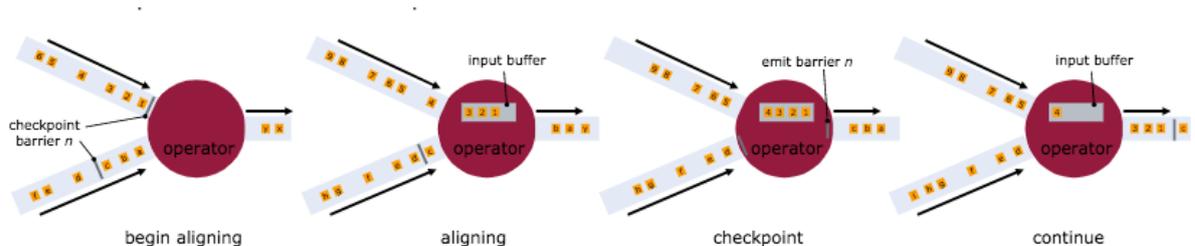
Each barrier carries the ID of the snapshot that covers the records before the barrier. Barriers are a lightweight mechanism. They do not interrupt the flow of the stream. Multiple barriers from different snapshots can be in the stream at the same time. This means that multiple snapshots may be created concurrently.

Figure 3-96: 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 S_n . 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.

Figure 3-97: Barrier mechanism



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.

3.4.2.4 Processing delay

Top 15 source subtasks with the longest processing delay

This metric describes the processing delays of each parallelism of a source.

3.4.2.5 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.

3.4.2.6 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.

3.4.2.7 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.

3.4.2.8 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.

3.4.2.9 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.

3.4.3 FailOver

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.

3.4.4 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. [Table 3-13: Parameter description](#) describes the parameters for the created checkpoints.

Table 3-13: 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. [Table 3-14: Parameter description](#) describes the parameters for the latest checkpoint.

Table 3-14: Parameter description

Parameter	Description
SubTask ID	The ID of the subtask.

Parameter	Description
State Size	The state size of the checkpoint.
Durations(ms)	The time that is spent on creating the checkpoint.

3.4.5 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.

3.4.6 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.

3.4.7 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.

3.4.8 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

On this tab page, you can view the basic running information of the current job.

***Table 3-15: Job properties* describes the basic job properties that are displayed on this tab page.**

Table 3-15: 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.

No.	Field and Description
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.

3.4.9 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 user-defined 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.



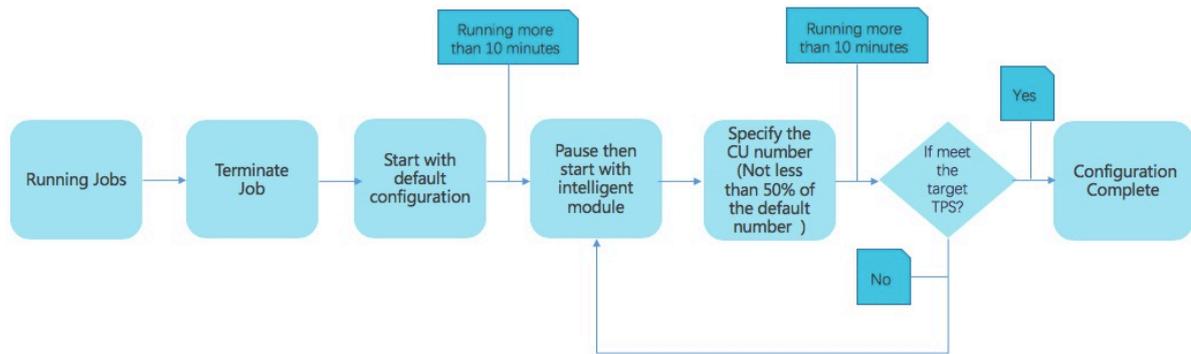
Note:

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.

Figure 3-98: Procedure



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 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

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.
      "targetResourceUnits": 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.IllegalStateException: 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(AbstractStreamOperator.java:275)
    at org.apache.flink.streaming.runtime.tasks.StreamTask.initialize
Operators(StreamTask.java:870)
    at org.apache.flink.streaming.runtime.tasks.StreamTask.initialize
State(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.SerializationExcepti
on: Cannot serialize/deserialize the object.
    at com.alibaba.blink.contrib.streaming.state.AbstractRocksDBRawSe
condaryState.deserializeStateEntry(AbstractRocksDBRawSecondaryState.
java:167)
    at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRe
storeOperation.restoreRawStateData(RocksDBIncrementalRestoreOperation.
java:425)
    at com.alibaba.blink.contrib.streaming.state.RocksDBIncrementalRe
storeOperation.restore(RocksDBIncrementalRestoreOperation.java:119)
    at com.alibaba.blink.contrib.streaming.state.RocksDBKeyedStateBac
kend.restore(RocksDBKeyedStateBackend.java:216)
    at org.apache.flink.streaming.api.operators.AbstractStreamOperator
.createKeyedStateBackend(AbstractStreamOperator.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.deserializeStateEntry(AbstractRocksDBRawSecondaryState.java:162)
... 11 more

```

3.4.10 Improve performance by manual configuration

3.4.10.1 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.

3.4.10.2 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 to HEAD. The values of the 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.

3.4.10.3 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
```

3.4.10.4 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.

Table 3-16: Parameter description

Object	Parameter	Description	Value
DataHub source table	<code>batchReadSize</code>	The number of data records that are read at a time.	Optional. Default value: 10.

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.

3.4.10.5 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.

3.4.10.6 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.

3.5 Apsara Bigdata Manager (ABM)

3.5.1 Routine maintenance

3.5.1.1 Perform routine maintenance

You can perform routine maintenance on Apsara Bigdata Manager (ABM) by using Apsara Infrastructure Management Framework.

Inspections in Apsara Infrastructure Management Framework

1. Log on to the ABM console.
2. Click  in the upper-left corner, and then click TIANJI to log on to the Apsara Infrastructure Management Framework console.
3. Go to the Clusters page in the ABM console and view the container status. Make sure that all containers are in the final status.
4. Go to the Dashboard page in the ABM console and view alerts. Make sure that no alert is generated.

Metrics and alert handling

- Hardware resource monitoring

Disk alert: The system saves logs of recent 30 days and deletes earlier logs. If a disk alert is generated when surged logs occupy too much disk space, contact technical support.

- System exception

If a system exception is found during the inspection, handle the exception in the ABM console. If the message about the exception is unclear, contact technical support.

3.5.1.2 View the ABM operating status

ABM monitors its own health and operating metrics. You need to regularly handle ABM alerts and view ABM operating metrics to evaluate system downtime risks in the future.

View ABM operating metrics

In ABM, click O&M on the top and click Clusters. The Overview tab appears.

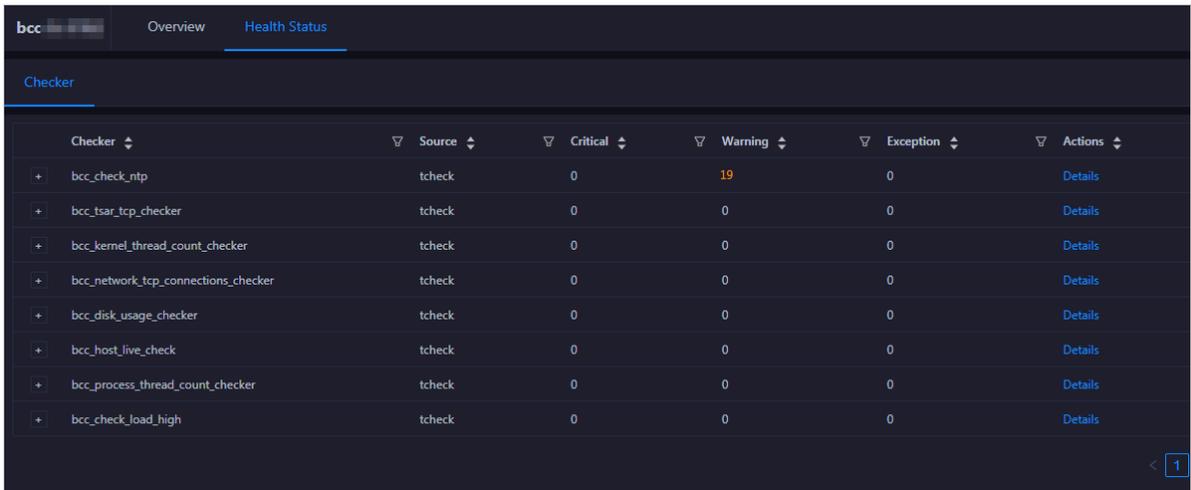


The Overview tab displays tendency charts for cluster metrics, including the CPU, memory, disk, load, package, TCP, and disk root directory usage. You need to regularly view and record these metrics to evaluate system downtime risks in the future.

Handle ABM alerts

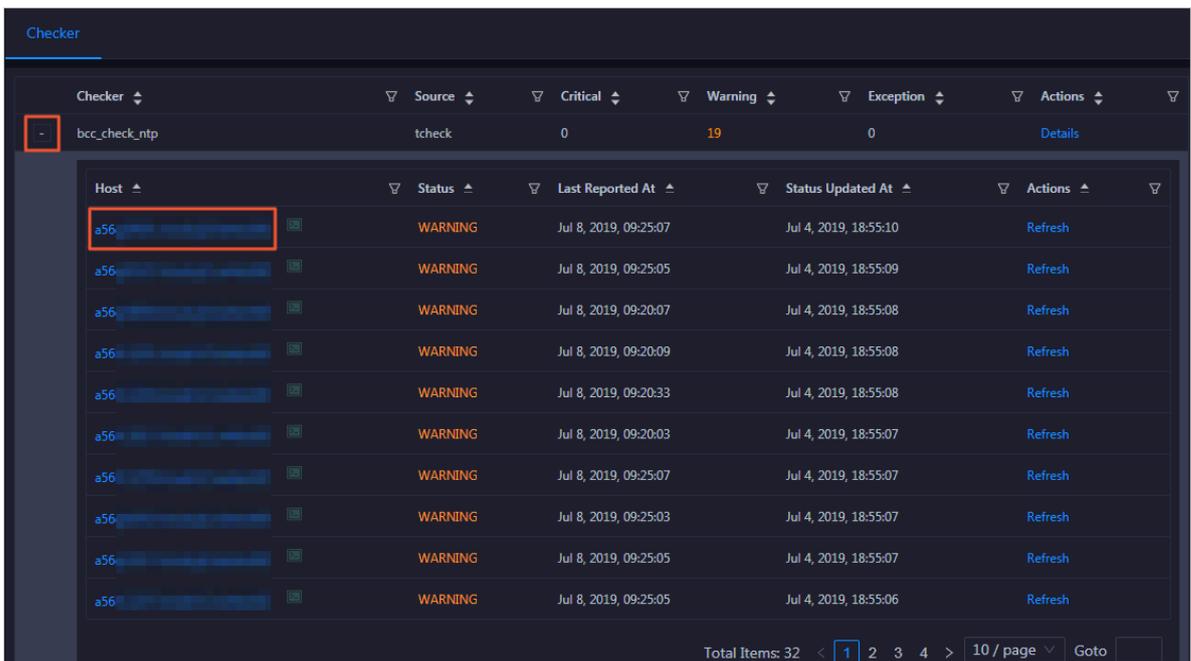
ABM cluster alerts are classified into Critical, Warning, and Exception alerts. You need to handle these alerts in time, especially Critical and Warning alerts.

1. On the Clusters page, click the Health Status tab.

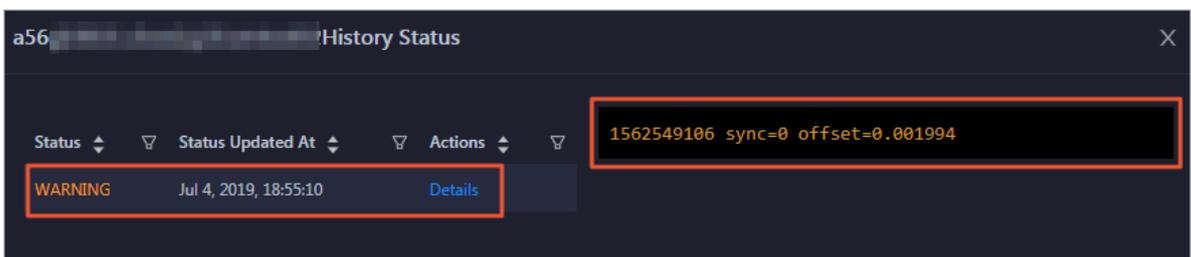


The Health Status tab displays all check items and the alerts that were generated during the check.

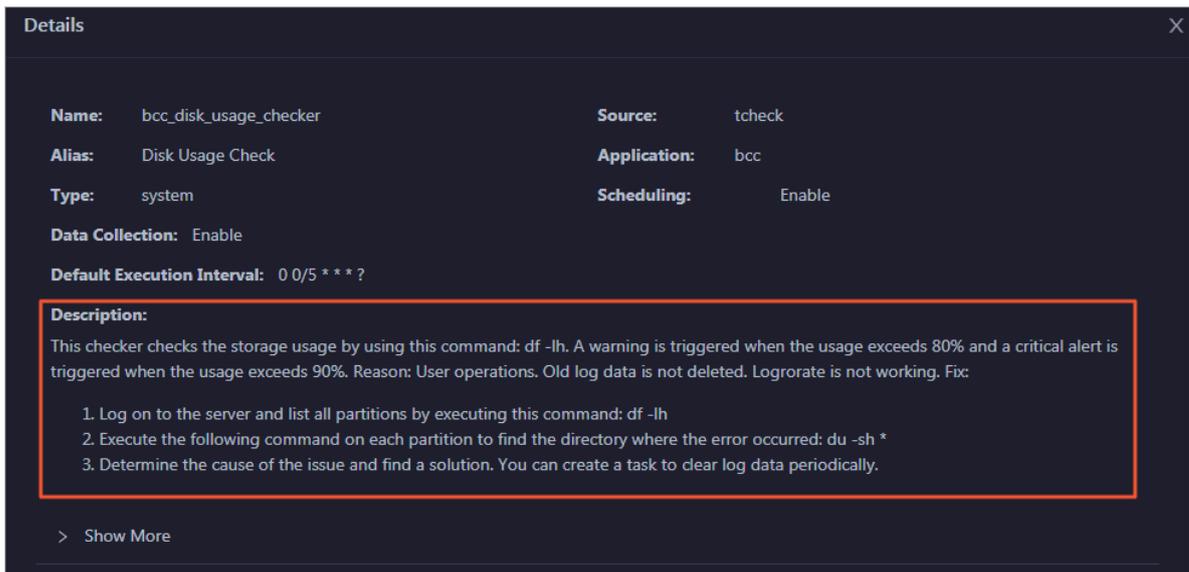
2. Click the Fold icon for a check item with alerts. All hosts on which the check item was performed appear.



3. Click a host. In the dialog box that appears, click Details for an alert. The alert cause appears on the right.



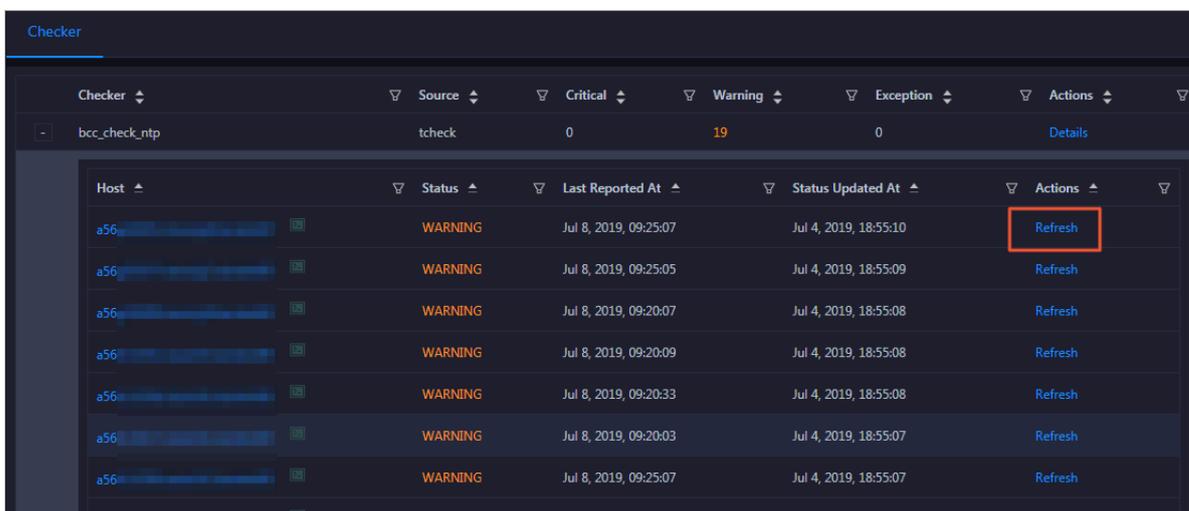
- Click Details for a check item with an alert and view the fix method for the alert in the dialog box that appears.



- Handle the alert based on the fix method.

You may need to log on to the host when handling the alert. For more information, see [Log on to a host](#).

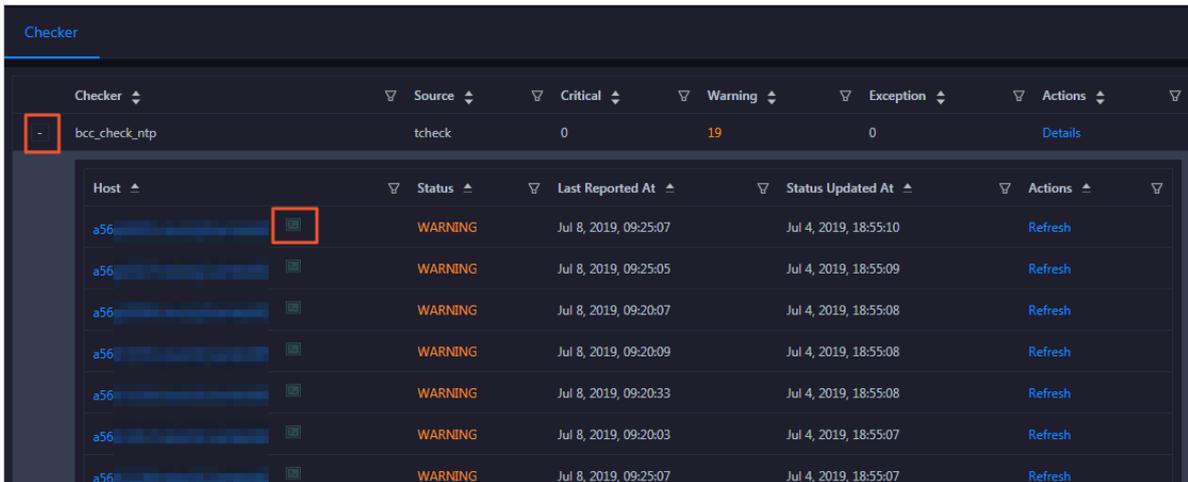
- After the alert is handled, click Refresh for the host to perform the check again in real time. In this way, you can check whether the alert is cleared.



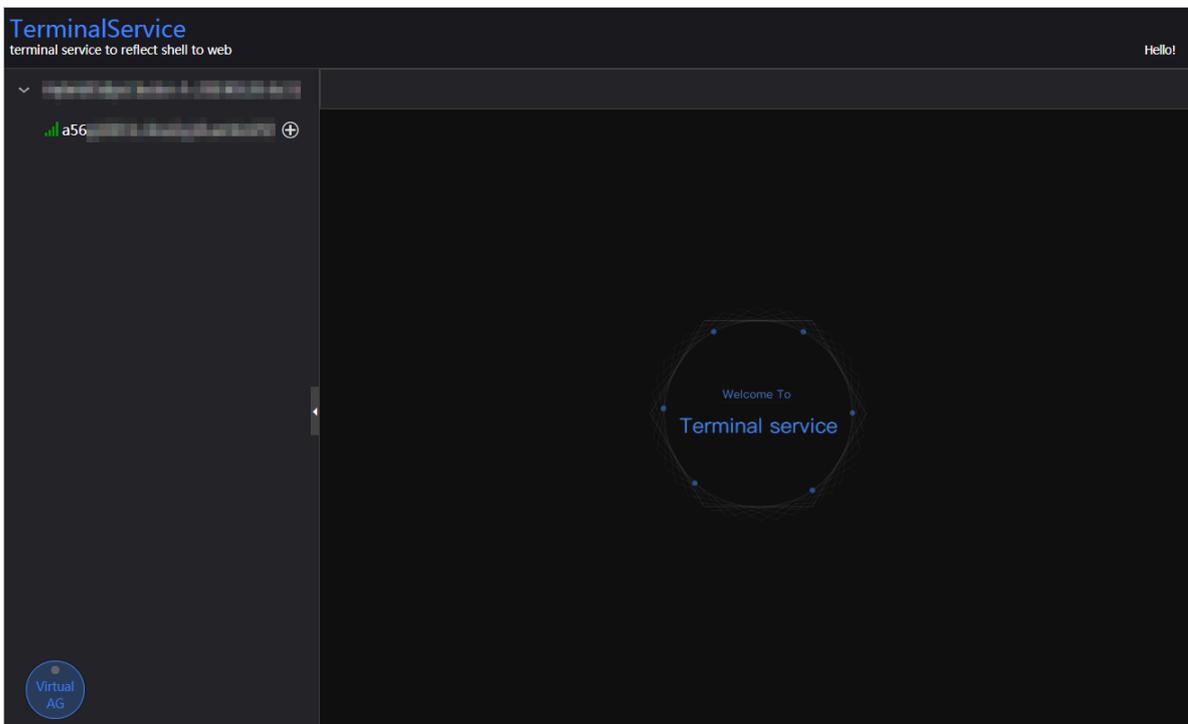
Log on to a host

You may need to log on to a host to handle alerts or other issues that occurred on the host.

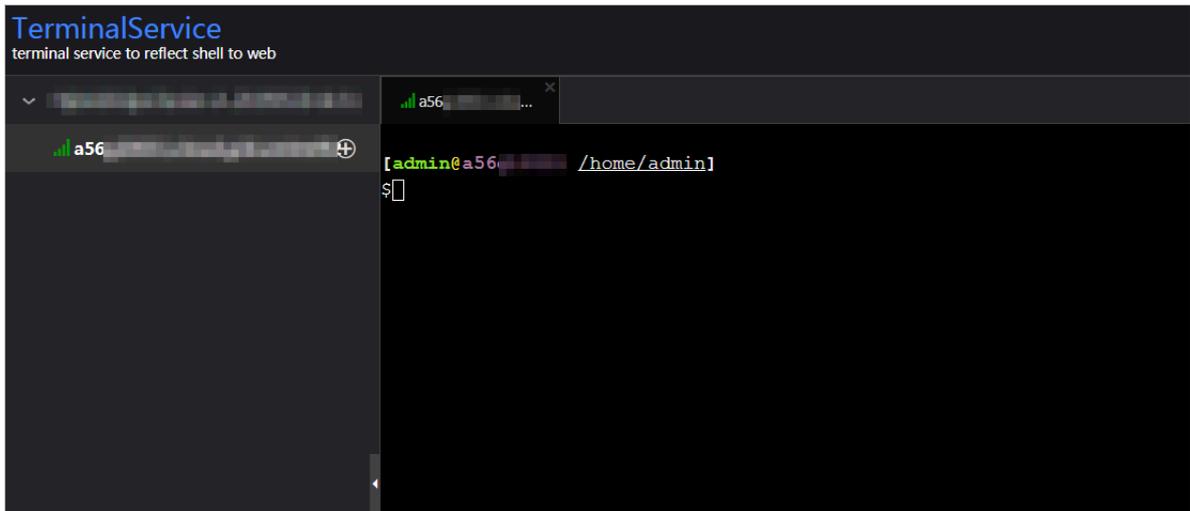
1. On the Health Status tab, click the Fold icon for a check item.



2. Click the Logon icon for a host. The TerminalService page appears.



3. On the TerminalService page, select the host on the left to log on to it.



3.5.1.3 Troubleshooting

Common failures

- **Logon failure**

If you failed to log on to ABM, clear the cache and cookies in your web browser, and then try again.

Based on the logon failure message that appears, check whether the following issues exist:

- The password that you entered is incorrect.
- Your account has been locked.
- Your account has been disabled.

- **Other failures**

Contact technical support.

3.5.2 Backup and restore

Back up data

ABM uses a high-availability database. You do not need to manually back up data. To obtain full backup data, contact technical support.

Restore data

You do not need to restore data for ABM.

3.6 Quick BI

3.6.1 Introduction to O&M and tools

3.6.1.1 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.

3.6.1.2 Troubleshoot Quick BI issues by using the Apsara Infrastructure Management Framework

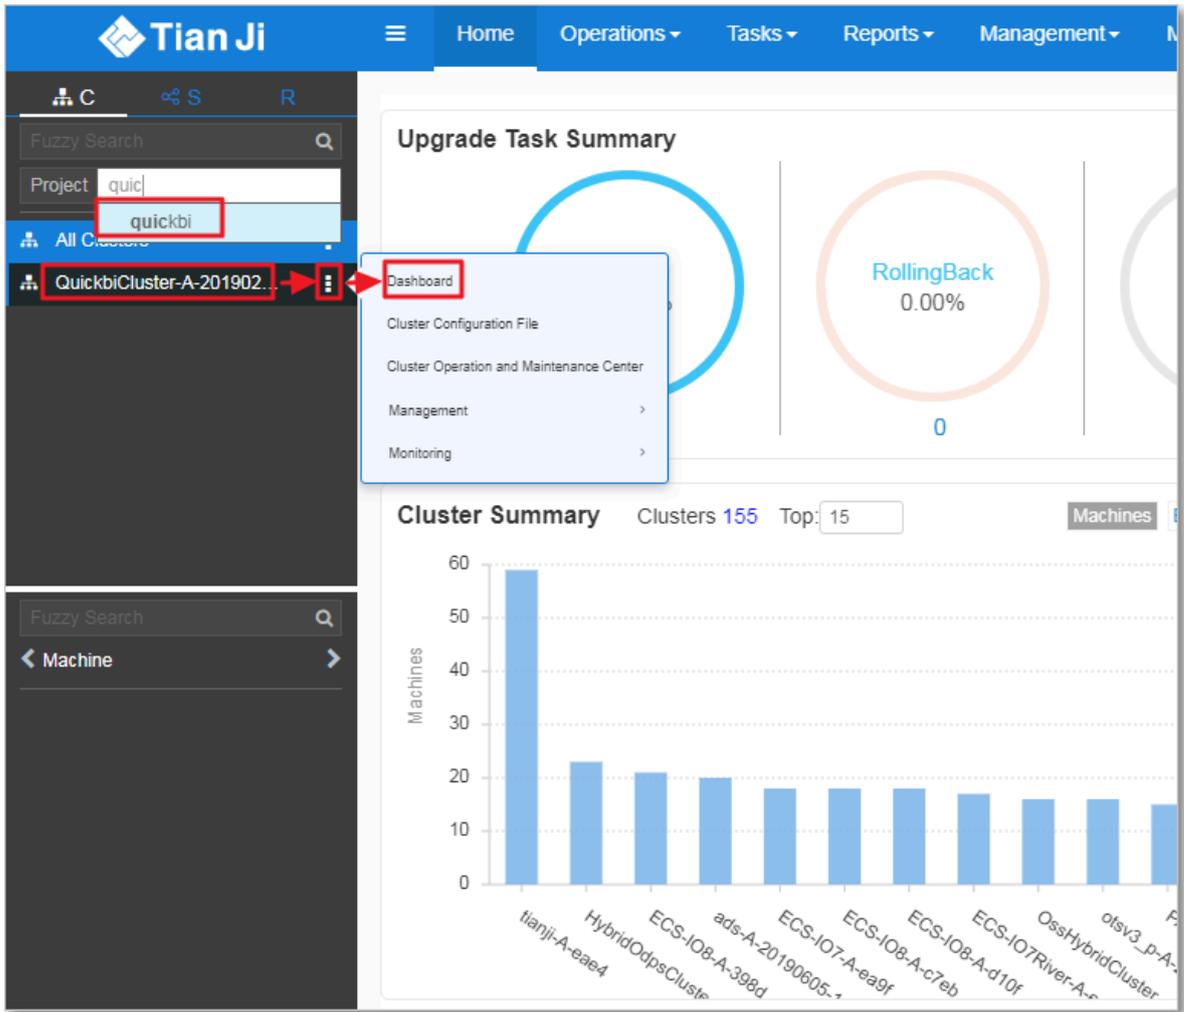
The Apsara Infrastructure Management Framework is a tool that allows you to perform O&M tasks on Quick BI. You can use the Apsara Infrastructure Management Framework to troubleshoot the service unavailability issue of Quick BI.

Prerequisites

Log on to the Apsara Infrastructure Management Framework.

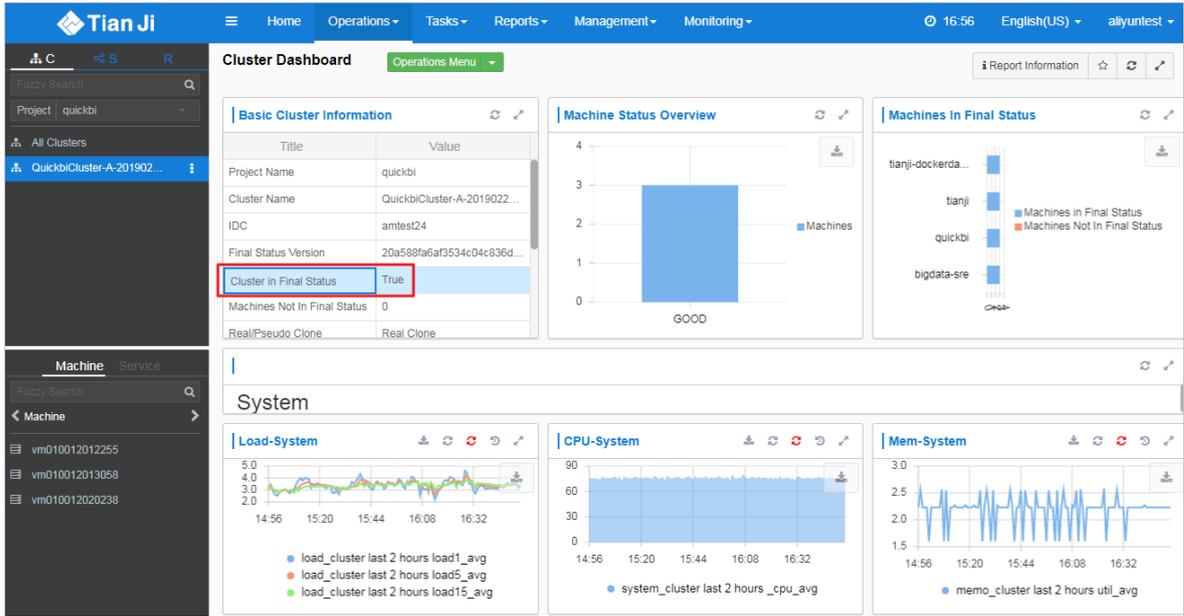
Procedure

1. Find the Quick BI project in the Apsara Infrastructure Management Framework.

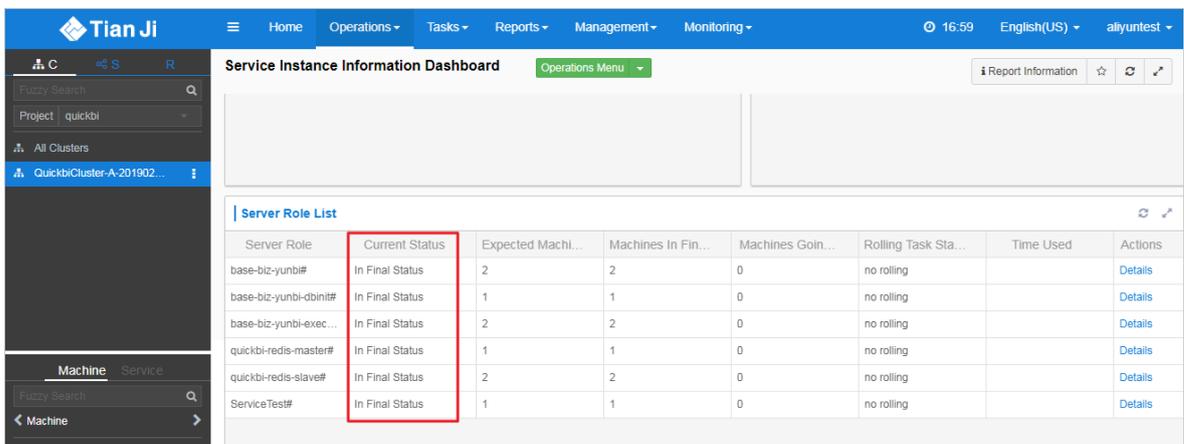
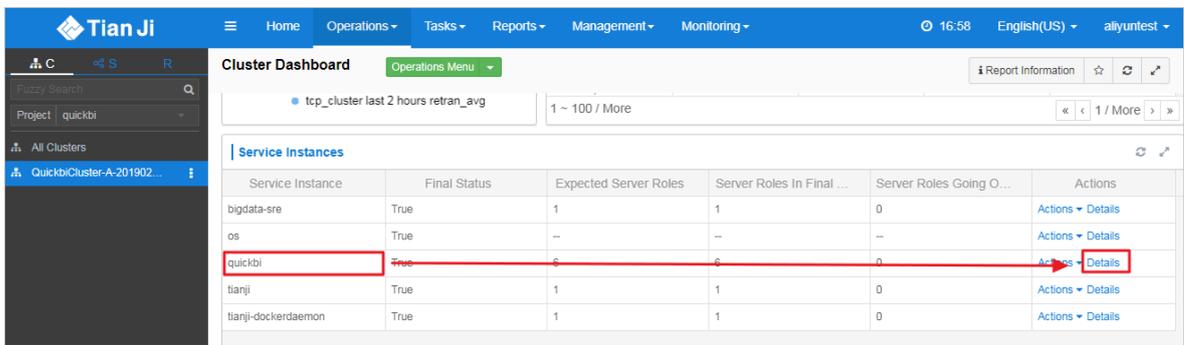


2. On the Dashboard page, view the cluster status of Quick BI. Check whether the Quick BI cluster is at the desired state. If the cluster is at the desired state, the

system works as expected. If the cluster is not at the desired state, go to the next step.



3. On the Dashboard page, find the Service Instances section, and view the service instance details of Quick BI.

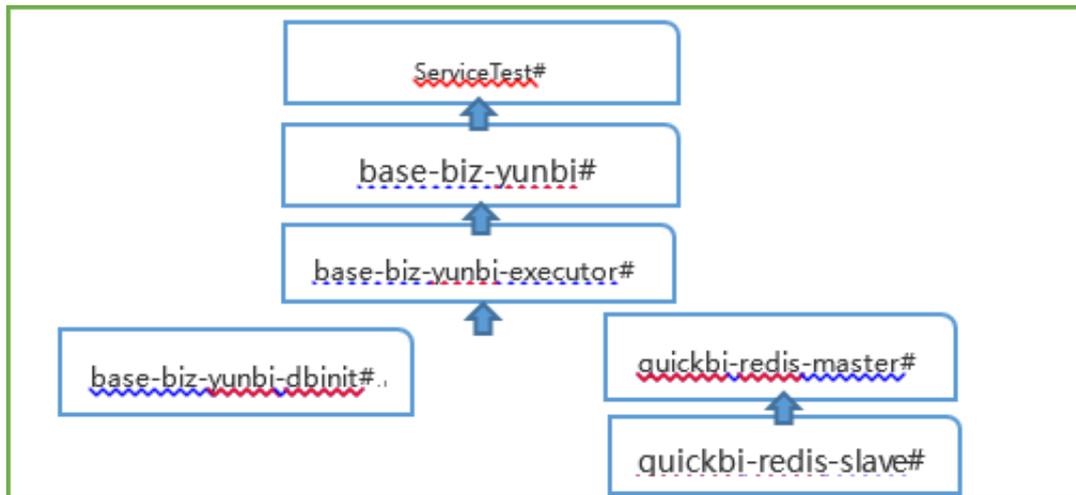


4. If a service instance is not at the desired state, you need to follow these steps to troubleshoot the issue.

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 *Figure 3-99: Relationship between Quick BI service roles* shows the relationship between Quick BI service roles.

Figure 3-99: 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.

3.6.2 Routine maintenance

3.6.2.1 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.

Quick BI O&M components, related service roles, and the description of 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 performance.
	quickbi-redis-slave#	
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.
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.

3.6.2.2 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-biz-yunbi-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.

3.6.2.3 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.

3.6.2.4 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.

3.6.2.5 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.

3.6.2.6 Automated testing components

This topic describes how to detect and troubleshoot issues when you perform container monitoring on automated testing components.

Container monitoring

Check whether the ServiceTest# service role is at the desired state.

A service unavailability issue is indicated if the check result shows that the service is at the desired state. Cause:

- **Error message:** The service is unavailable. **Symptom:** You cannot visit <https://container/checkpreload.htm> and the Quick BI console.



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.

- **The service is available but the check result shows errors.** **Symptom:** You can log on to the Quick BI console and search data. However, the Monitoring page of Apsara Infrastructure Management Framework shows a logon error. This issue occurs when you use the default base_admin@aliyun.com account to run test cases, the password for the account is changed due to re-deployment of MaxCompute (previously known as ODPF). In this case, the value of the `${service:odps-service-computer:base_meta_account.password}` environment variable is incorrect. Therefore, you cannot log on to the Quick BI console. You can view the error message provided in the Description column.

Solutions:

- **Solution for the issue that the service is unavailable. If other service roles are not at the desired state, you need to follow the predefined process to handle this issue.**
- **Solution for the issue that the service is available but the check result shows errors. If the password for the base_admin@aliyun.com account is changed, you need to restart the ServiceTest service role.**

Periodical detection

You can run test cases at regular intervals to check the availability of Quick BI. A service is available if the linked service role is at the desired state. Otherwise, the service is unavailable. The detection interval is 30 minutes.

3.7 Graph Analytics

3.7.1 Operations and maintenance tools and logon methods

3.7.1.1 Log on to Apsara BigData Manager

This topic describes how to log on to Apsara BigData Manager.

Prerequisites

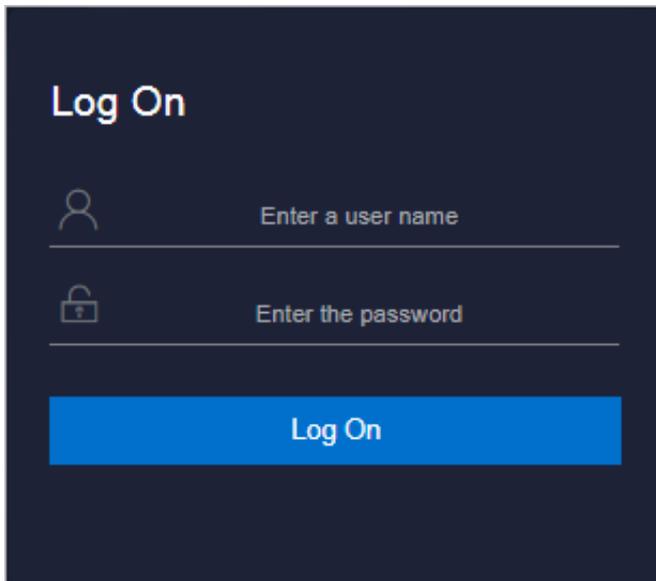
- **ASO access address in the format of `http://region-id.aso.intranet-domain-id.com`.**
- **Google Chrome browser (recommended).**

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.

Figure 3-100: Log on to ASO

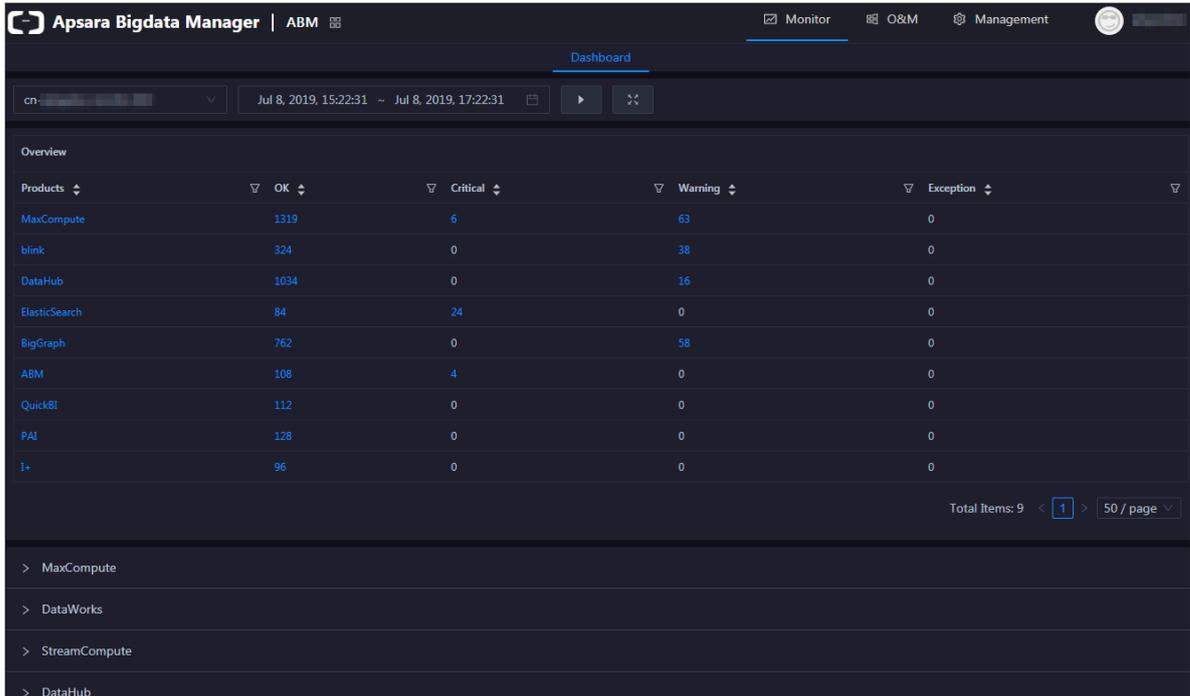


Note:

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.

- In the left-side navigation pane, choose **Products > Apsara BigData Manager** to go to the homepage of Apsara BigData Manager.



- On the homepage of Apsara BigData Manager, click the **I+** icon, and then select **I+** to go to the page.



3.7.1.2 Log on to Apsara Infrastructure Management Framework

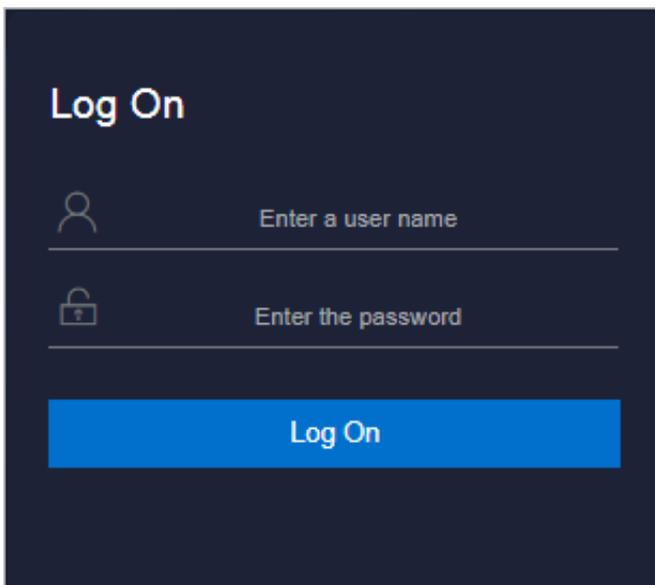
Prerequisites

- 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.

Figure 3-101: Log on to ASO



Note:

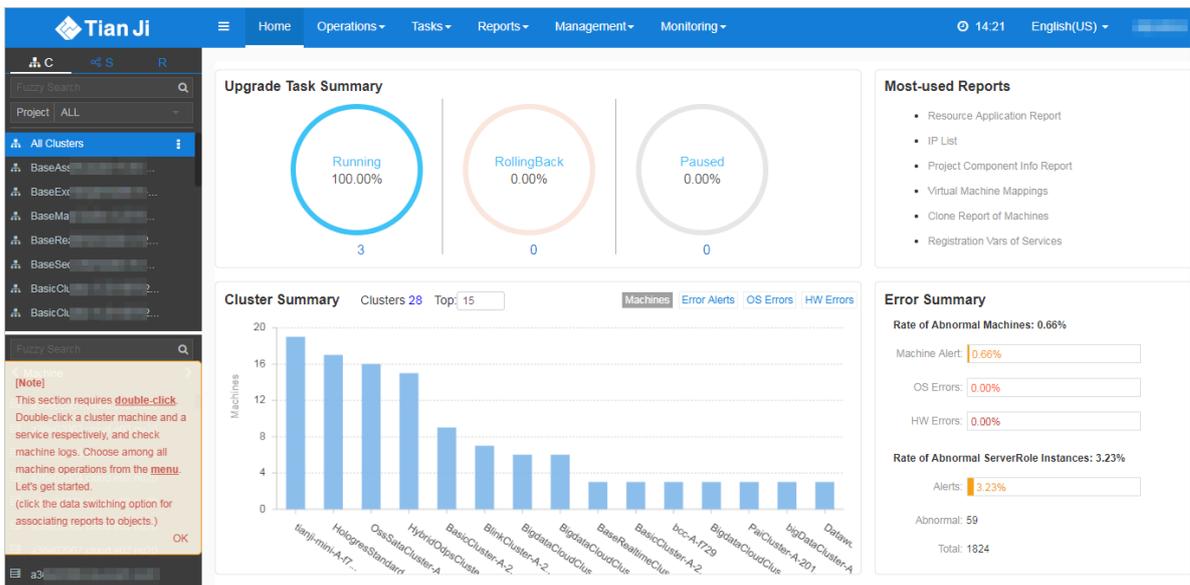
You can select the language from the drop-down list in the upper-right corner to change the language of ASO.

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.
5. In the left-side navigation pane, choose Products > Apsara Infrastructure Management Framework to go to the homepage of Apsara Infrastructure Management Framework.

Figure 3-102: Homepage of Apsara Infrastructure Management Framework



3.7.1.3 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.

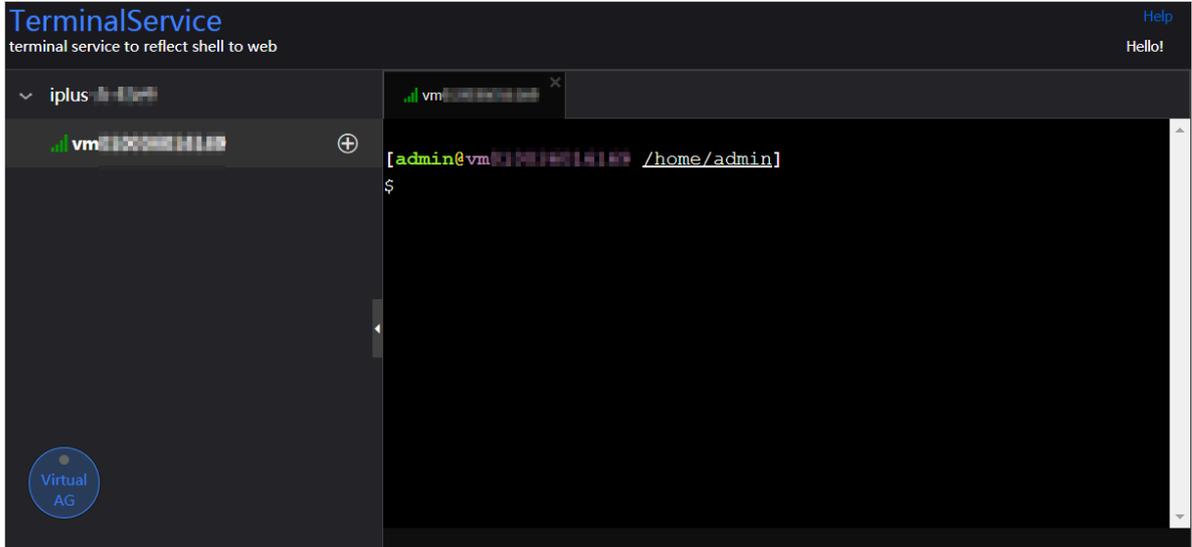
Procedure

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 `vmXXXXXXXXXXXXXX` 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.

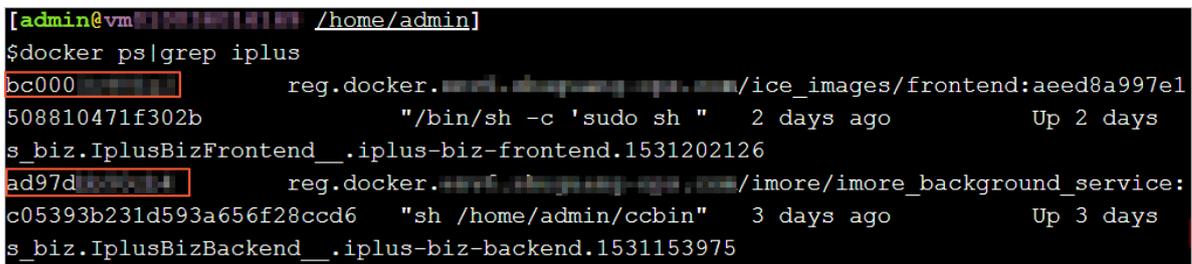
Figure 3-103: TerminalService page



The left-side navigation pane on the TerminalService page displays the virtual machine selected by you (vmXXXXXXXXXXXXX).

4. In the left-side navigation pane on the TerminalService page, click vmXXXXXXXXXXXXX, 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.

Figure 3-104: Query the docker ID



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`.

Figure 3-105: Log on to the docker container

```
[admin@vm010 169 /home/admin]
$docker exec -ti ad97d bash
[root@docker01 71 /home/admin]
```

7. The root account is used by default. You can use the `su - admin` command to switch to the admin account.

Figure 3-106: Switch to the admin account

```
[root@docker01 71 /home/admin]
#su - admin
[admin@docker01 71 /home/admin]
$
```

3.7.2 Operations and maintenance

3.7.2.1 Operations and maintenance based on BigData Manager

3.7.2.1.1 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

Your have obtained an Apsara BigData Manager account and the password with Graph Analytics O&M permissions.

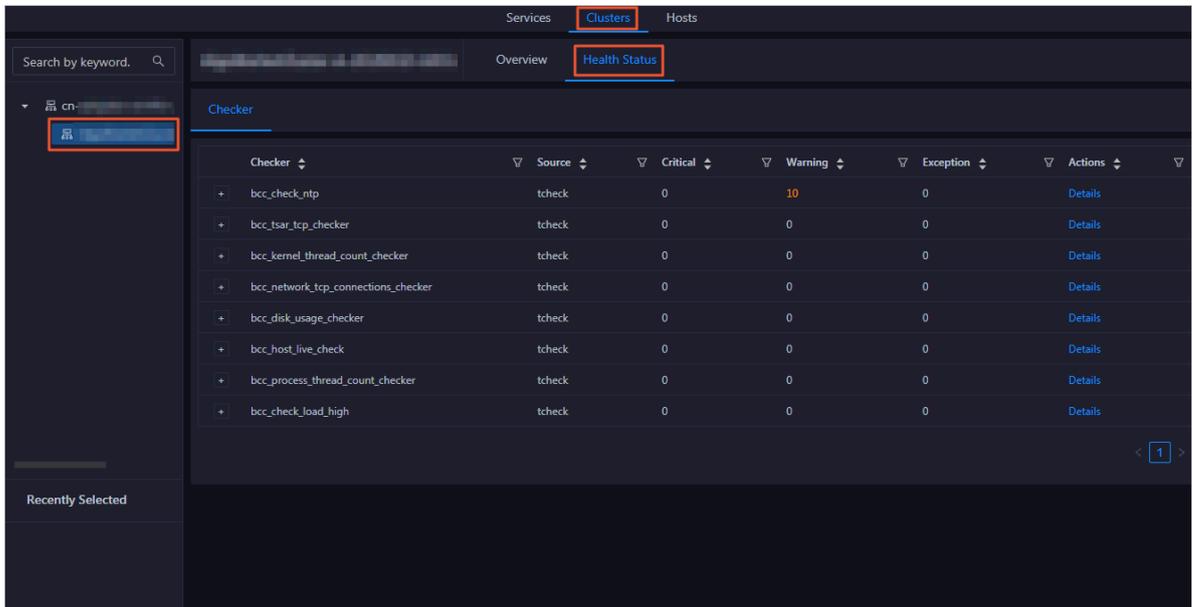
Step one: View cluster alerts

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.



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.

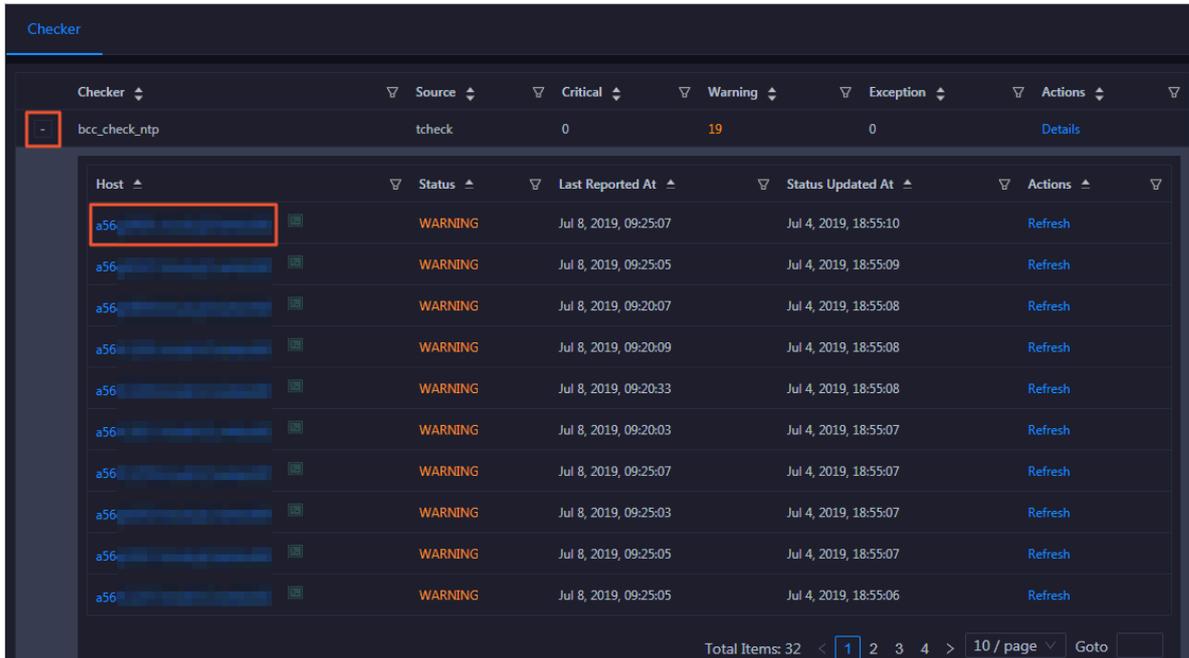


The Health Status tab page displays all health check items of the current cluster. You need to focus on the check items with Critical and Warning alerts.

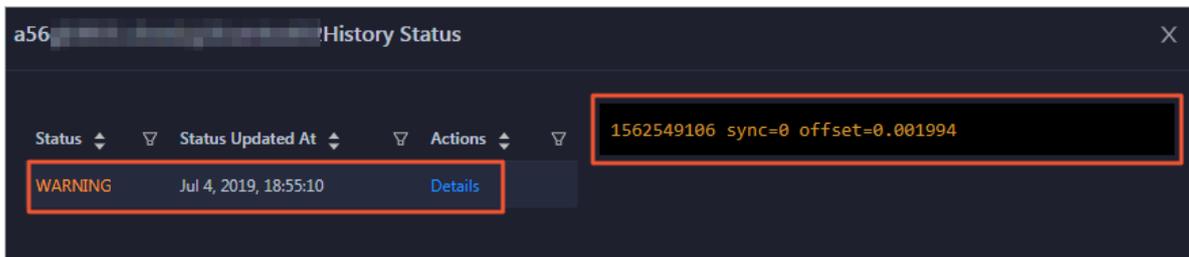
Step two: View hosts in the alert status and the alert causes

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.



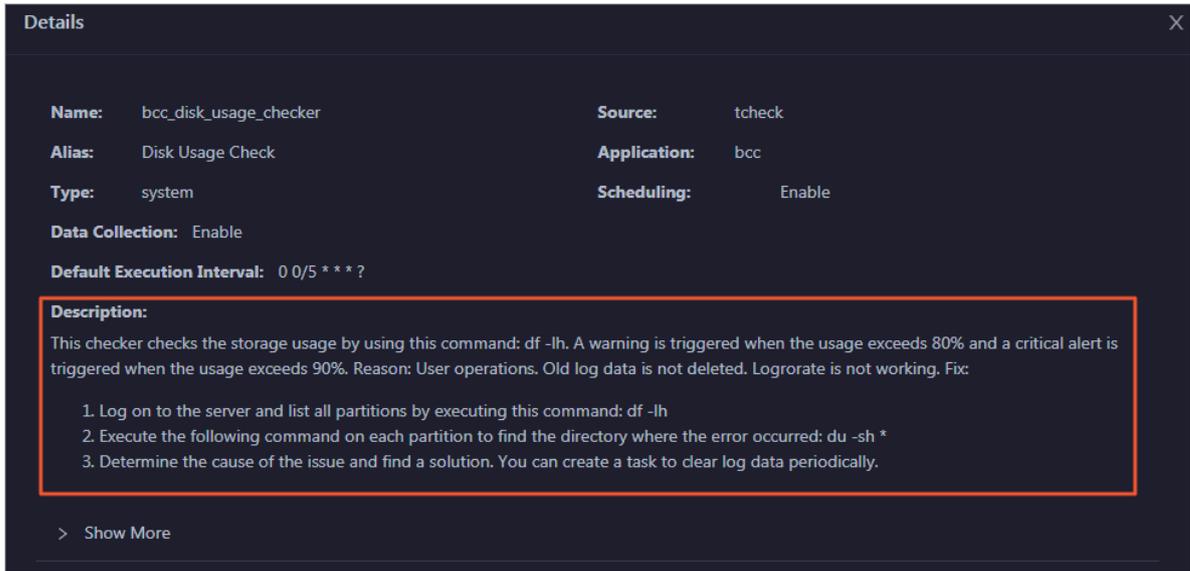
2. Click a host. In the dialog box that appears, click Details, and the cause of the alert appears on the right side.



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.

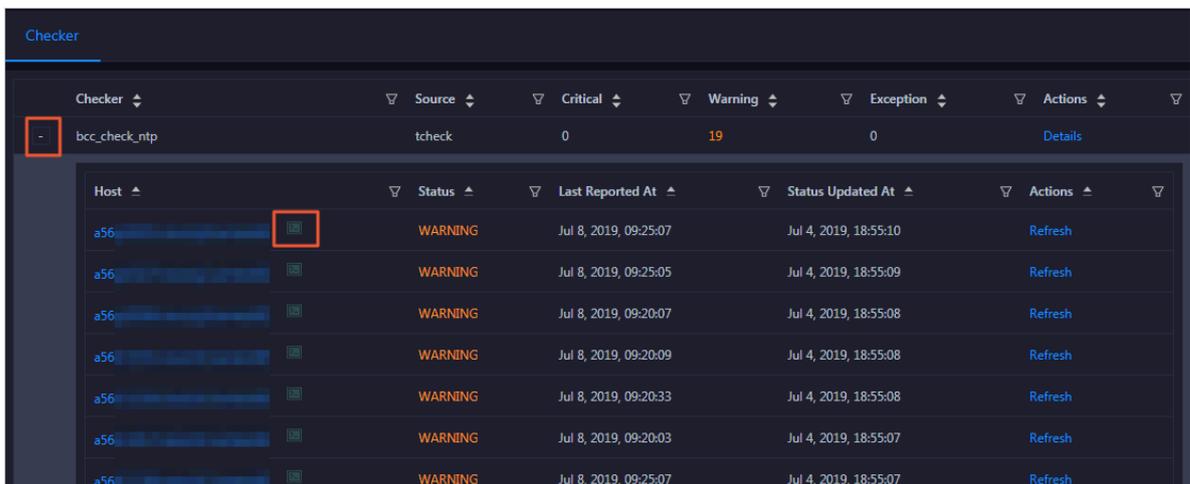


2. Handle the alerts based on the procedure described in Fix.

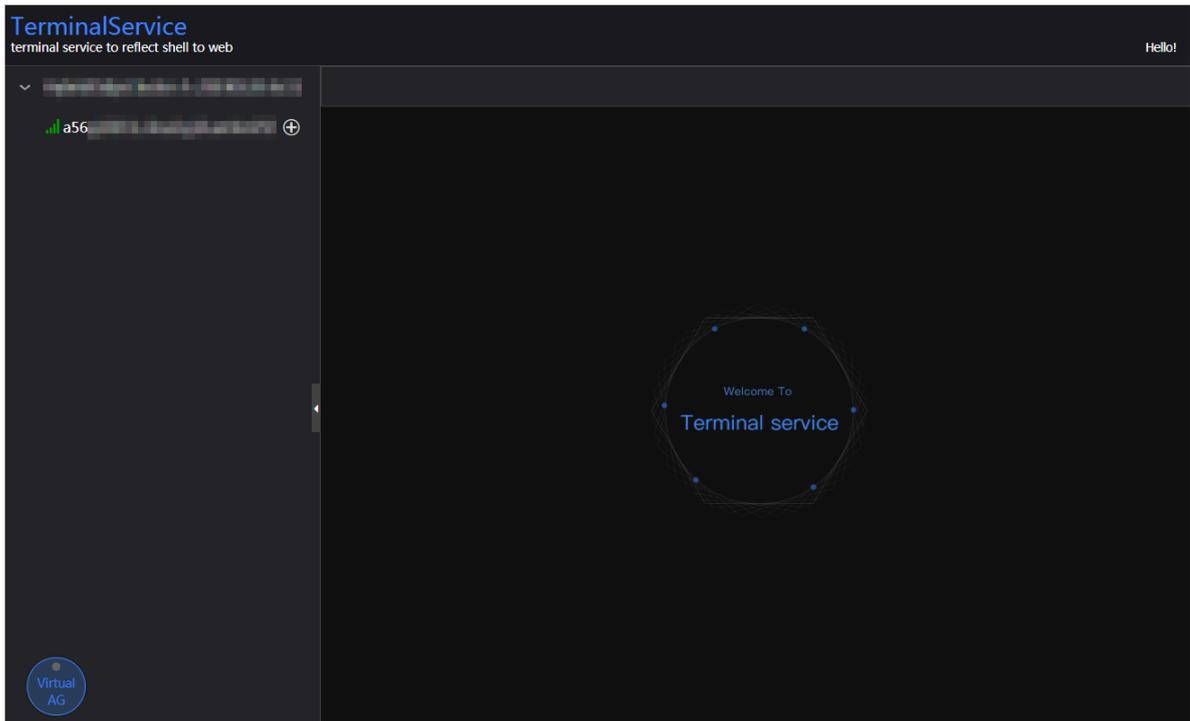
When you handle host alerts, you may need to log on to the host to perform related operations. For more information about how to log on to the host, see [Step four: Log on to a host](#).

Step four: Log on to a host

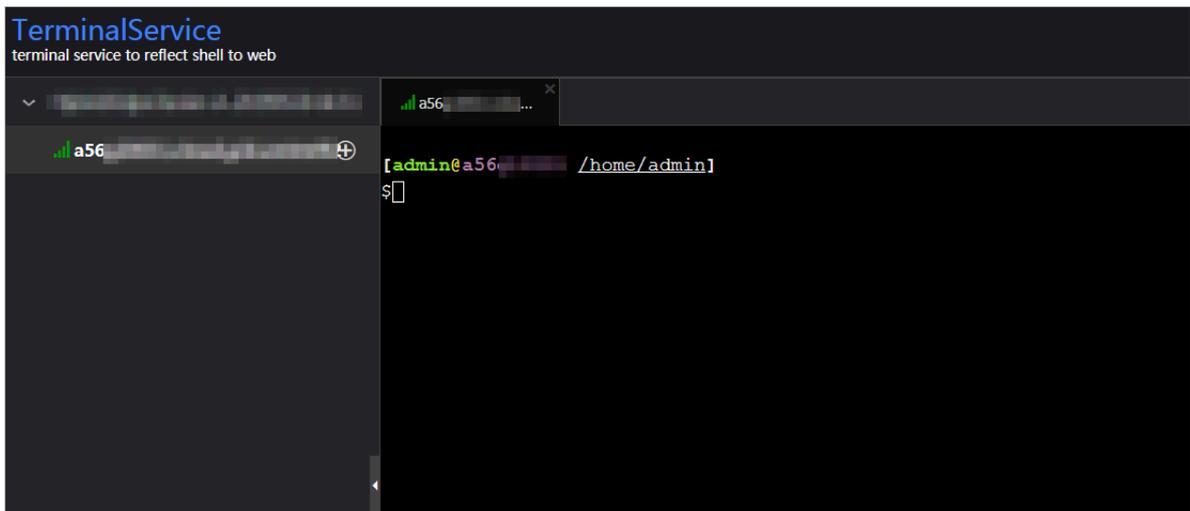
1. On the Health Status tab page, click the Plus sign (+) of a check item.



2. Click the Logon icon of a host. The TerminalService page occurs.



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.



3.7.2.1.2 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

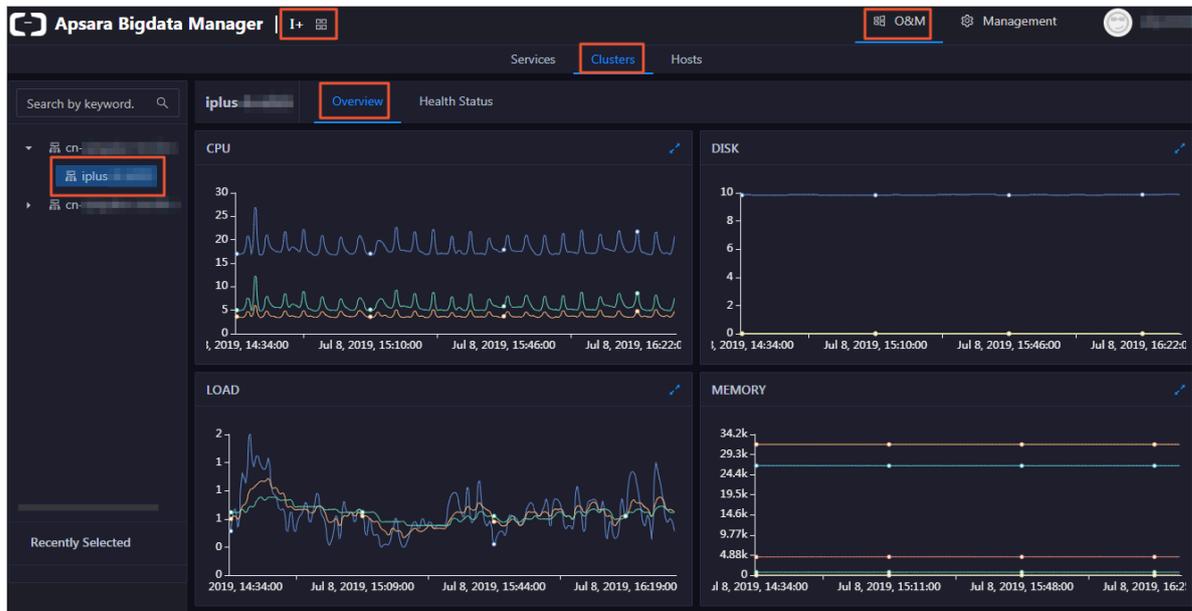
detected, the IT administrator must take immediate measures to identify the cause and fix the issue.

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.

3.7.2.1.3 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

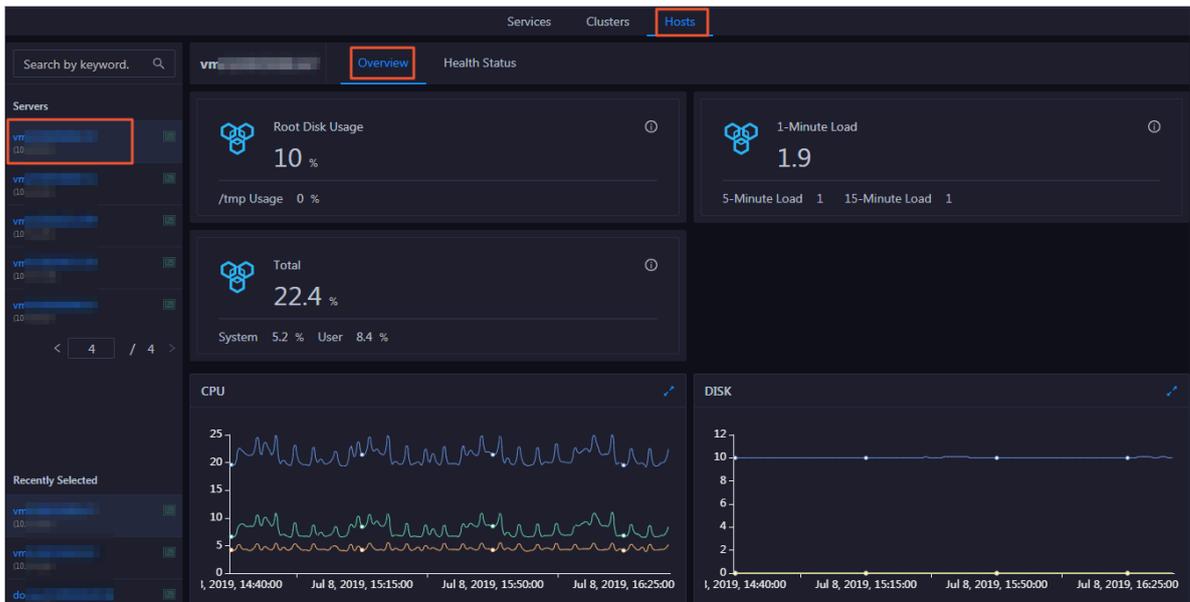
detected, the IT administrator must take immediate measures to identify the cause and fix the issue.

Prerequisites

You 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.



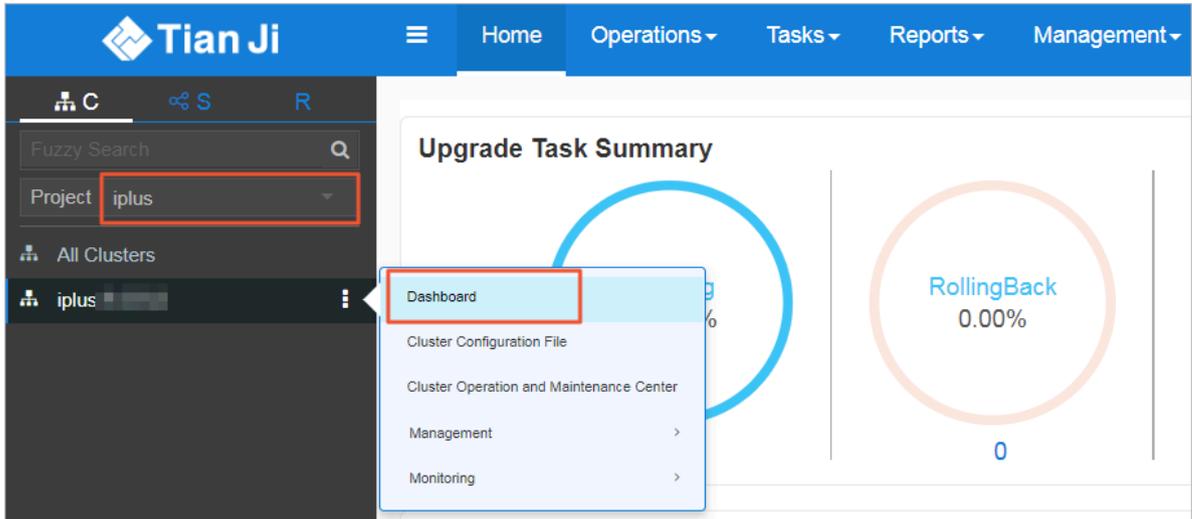
The Overview tab page displays the root disk usage, CPU usage, disk, memory, load, package, and TCP.

3.7.2.2 Operations and maintenance based on Apsara Infrastructure Management Framework

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

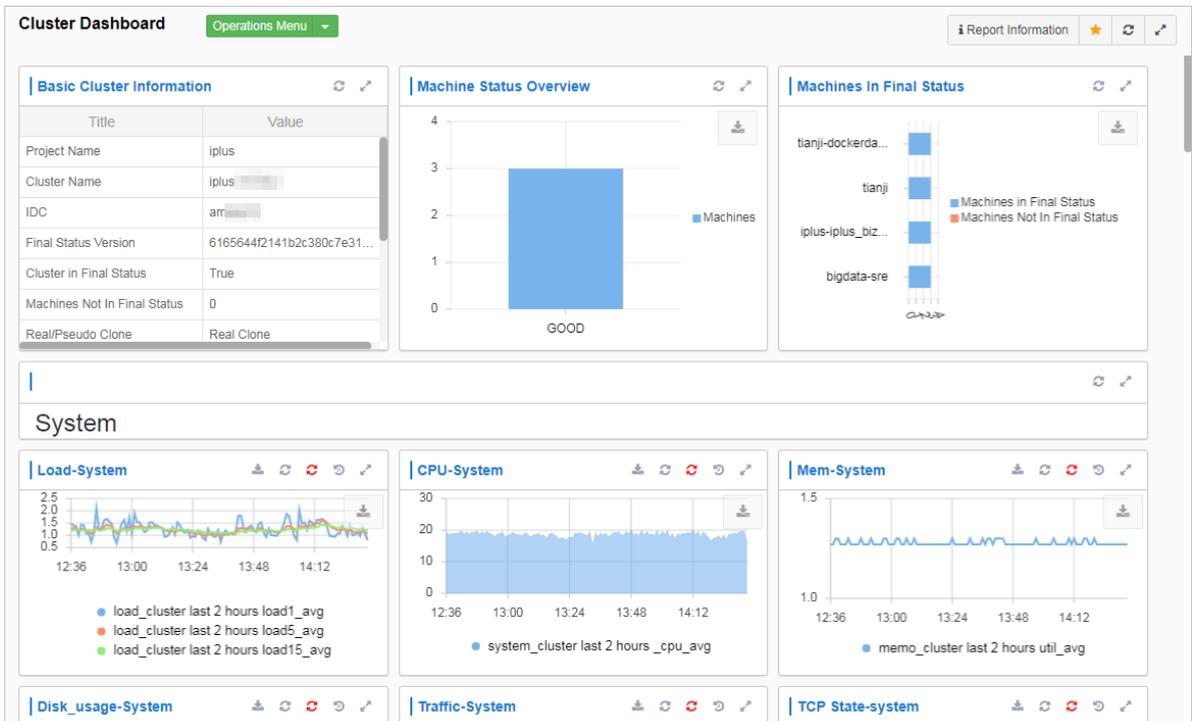
2. Enter **iplus** in the search box to search for the **iplus** cluster, as shown in *Figure 3-107: Search for the iplus cluster.*

Figure 3-107: 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 *Figure 3-108: Cluster Dashboard page.*

Figure 3-108: Cluster Dashboard page



4. In the Service Instances list, click Details for iplus-iplus_biz. The Service Instance Information Dashboard page appears, as shown in *Figure 3-109: Service instance list*.

Figure 3-109: Service instance list

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.

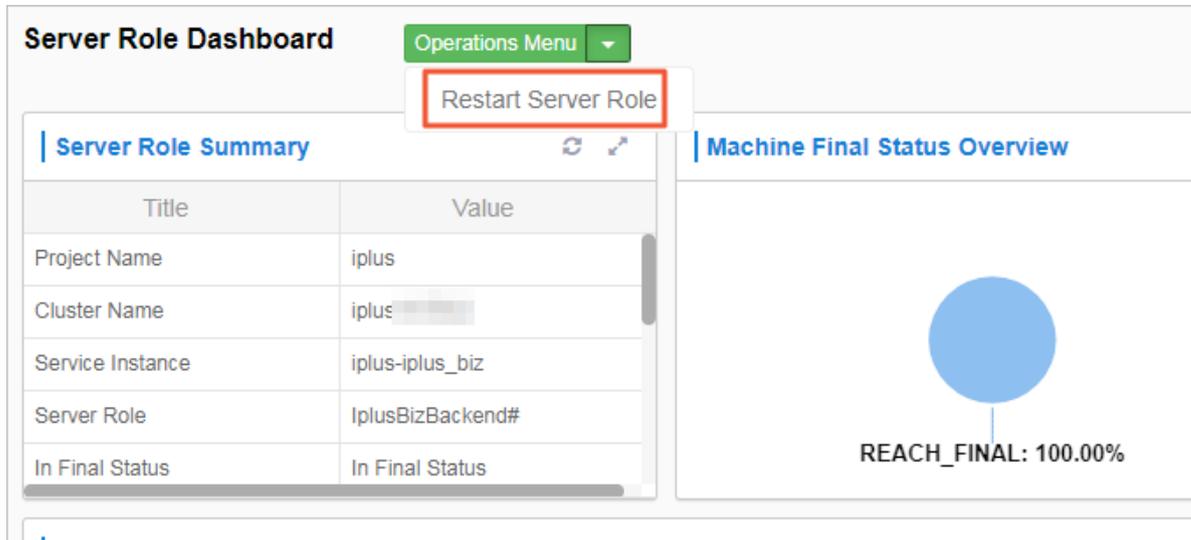
Other roles can be restarted in any order.

Figure 3-110: Server role list

Server Role	Current Status	Expected Machi...	Machines In Fin...	Machines Going...	Rolling Task Stat...	Time Used	Actions
IplusBizBackend#	In Final Status	3	3	0	no rolling		Details
IplusBizBackendContr...	In Final Status	2	2	0	no rolling		Details
IplusBizDbinit#	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](#).

Figure 3-111: Restart server roles



3.7.2.3 Operations and maintenance based on the Graph Analytics container

3.7.2.3.1 View instances

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 [Figure 3-112: View Java running instances](#) exists, Administration Console is in the normal status.

Figure 3-112: View Java running instances

```

sps -ef|grep java|grep iplus
admin 26378 1 0 Jul05 ? 00:24:36 java -server -xms1800m -mx1800m -xmn600m -xss256k -XX:PermSize=512m -XX:MaxPermSize=512m -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/home/admin/1
logs -XX:+UseConcMarkSweepGC -XX:+UseParNewGC -XX:+CMSFullGCsBeforeCompaction=5 -XX:+UseCMSCompactFullCollection -XX:+CMSClassLoadingEnabled -XX:+DisableExplicitGC -verbose:gc -XX:PrintGCDetails -X
X:+PrintGCInstants -DFILE_ENCODING=UTF-8 -jar /home/admin/iplus_pack/iplus-control.war --spring.config.location=/home/admin/iplus_pack/config/application-control.yml
admin 27322 1 0 Jul05 ? 00:23:50 java -server -xms5000m -mx5000m -xmn1024m -xss256k -XX:PermSize=512m -XX:MaxPermSize=512m -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/home/admin/
logs -XX:+UseConcMarkSweepGC -XX:+UseParNewGC -XX:+CMSFullGCsBeforeCompaction=5 -XX:+UseCMSCompactFullCollection -XX:+CMSClassLoadingEnabled -XX:+DisableExplicitGC -verbose:gc -XX:PrintGCInstants
s -XX:+PrintGCDetails -XX:+PrintHeapAtGC -Xloggc:/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.war --spring.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 *Figure 3-113: View node instances* exists, the node service of Graph Analytics is normal.

Figure 3-113: View node instances

```

$ps -ef|grep node
admin      7974      1    0 19:12 pts/0    00:00:00 node /home/admin/i3-admin/target/i3-admin/admin-patch.js --harmony
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.js --harmony 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.js --harmony 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.js --harmony 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.js --harmony undefined
admin      14876     1    0 Aug16 ?        00:00:00 node /home/admin/i3-web/target/i3-web/dispatch.js --harmony
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.

3.7.2.3.2 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.

3.7.2.3.3 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`.**

Figure 3-114: View database information

```
datasource:
  url: jdbc:mysql://iplus-meta.mysql.aliyuncs.com:3177/iplus_meta?useUn
eSSL=true
  username: iplus_meta
  password: iplus_meta
  driver: com.mysql.jdbc.Driver
  type: com.alibaba.druid.pool.DruidDataSource
  druid:
    max-active: 50
    initial-size: 1
    min-idle: 3
    max-wait: 60000
    time-between-eviction-runs-millis: 60000
    min-evictable-idle-time-millis: 300000
    test-while-idle: true
    test-on-borrow: false
    test-on-return: false
```

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.manual%' and (gmt_create < now() and
gmt_create > date_sub(now(), interval 1 hour) );
```

3.7.2.3.4 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.

3.7.2.3.5 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`

3.7.3 Security maintenance

3.7.3.1 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.

3.7.3.2 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.

3.7.4 Troubleshooting

3.7.4.1 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.

3.7.4.2 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.

3.7.4.3 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.

3.7.4.4 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.**

3.8 Machine Learning Platform for AI

3.8.1 Query server and application information

3.8.1.1 Apsara Stack Machine Learning Platform for AI

3.8.1.1.1 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.
 - a) Find the service instance and click Details. The instance detail page appears.
 - b) Find the role list and click Details. The role detail page appears.
 - c) 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.

3.8.1.1.2 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.

3.8.1.1.3 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/classes/application.yml` file.**



Note:

In the preceding file path, `{app}` indicates the component name, such as `pai-dms`.

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.**

- a) **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.**
- b) **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
```

3.8.1.1.4 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
```

3.8.1.2 Online model service

3.8.1.2.1 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.**

3.8.1.2.2 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
```

3.8.1.2.3 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}
```

3.8.1.3 GPU cluster and task information

3.8.1.3.1 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 ttrtl` 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.

3.8.1.3.2 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 al` command.

3.8.2 Maintenance and troubleshooting

3.8.2.1 Machine Learning Platform for AI maintenance

3.8.2.1.1 Run ServiceTest

After `ServiceTest` is run, the automated test case is executed.

1. 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.

Figure 3-115: ServiceTest instance

pai	Final	21 Hours 19 Minutes	Cluster: 4 / 4	Service: 18 / 18	Role: 23 / 23	Total: 21	Done: 21	0	0	✖
	Final	21 Hours 20 Minutes	AlgoMarketClust...	bigdata-sre	PaiAlgoinit#			0	0	
	Final	21 Hours 20 Minutes	AllinkCluster-A-2...	os	PaiDbinit#			0	0	
	Final	21 Hours 20 Minutes	EASCluster-A-20...	pai-pai_service	PaiDmscloud#			0	0	
	Final	21 Hours 20 Minutes	PaiCluster-A-20...	tianji	PaiFront#			0	0	✖
	Final	1 Hour 7 Minutes		tianji-dockerdae...	PaiMemcached#			0	0	✖
	Final	21 Hours 20 Minutes			ServiceTest#			0	0	✖
	Final	21 Hours 18 Minutes						0	0	✖
	Final	11 Hours 48 Minutes						0	0	

6. Run the `sudo docker restart e90f70353031` command to restart the ServiceTest service, as shown in the following figure.

Figure 3-116: Restart the ServiceTest service

```

sudo docker ps -a
CONTAINER ID        IMAGE               PORTS              NAMES
STATUS
e90f70353031      inc.com/idst-pai/pai-web-test:db13d8a23beebc5495751d86d856ef51  inc.com/idst-pai-pai-web-test-1bc97  "sh /usr/local/smokin"  10 days ago
Exited (0) About an hour ago
pai-pai_service.ServiceTest_..._service_tes

```

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.

Figure 3-117: Testing results

```

[admin@vm010036032130 /home/admin]
$ sudo docker restart e90f70353031
e90f70353031

```

- **PASS:** The algorithm is running properly.
- **SKIP or FAIL:** The algorithm fails.

3.8.2.1.2 Common faults and solutions

3.8.2.1.2.1 Maintenance commands

`nc`, `telnet`, `curl`, `ping`, `mysql`

`docker images` : shows all images on a server.

`docker ps`: shows the running images on a server.

`docker exec -ti containerID /bin/bash`

`docker log containerID`: shows the container log.

`curl http://localhost/status.taobao`: determines whether the SpringBoot service is started.

3.8.2.1.2.2 *pai.xx.xx access failures*

Procedure

1. Run the `ping pai.xx.xx` command to check whether the domain name has been translated to the corresponding VIP.

If the domain name cannot be resolved properly, contact the on-site engineer to check the network configurations.

2. Run the `curl http://ip/status.taobao` command to check whether all service modules are running normally.

If the `status.taobao` module fails the check, perform the following operations:

- a. Log on to the server to check whether the container is active.
- b. Go to the container and run the following command to check whether the service process is active:

```
ps -lef | grep java
```

- c. View the `/home/admin/{app}/logs/err_pai-dms.log` file to locate causes, such as dependent tenant service request timeout, dependent OCS timeout, and database connection exceptions.

We recommend that you view the log after checking all items in the checklist to verify whether the malfunction was not caused by a component exception.

3. Verify whether ApsaraDB for RDS is accessible.

- a) Run the following command to check whether the port is active:

```
nc -v -z $rds_host $port
```

- b) Run the following command to check whether the database is accessible:

```
mysql -h$Host -P$Port -u$user -p$password
```

4. Verify whether the caching service is functioning properly.

Run the following command to check whether port 11211 is active:

```
nc -v -z $ocs_host 11211
```

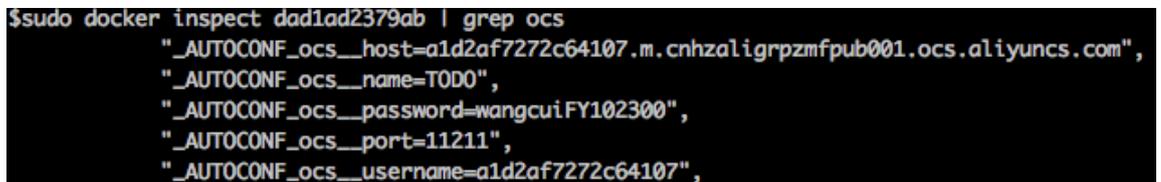
Search for `ocs_host` as follows:

a. Search for the `dmscloud` instance, as shown in the following figure.



```
[admin@vm010036008128 /home/admin]
$ sudo docker ps
CONTAINER ID        IMAGE                                     NAMES
STATUS            PORTS
b6ead0fa1d58      aliyun-inc.com/1dst-pai/dmscloud:      pai-pai_service.PaiDmscloud
Up 10 days         pai-pai_service.PaiDmscloud .pai_dmscloud.1519922511
```

b. Run the `sudo docker inspect b6ead0fa1d58 | grep ocs` command to view the `ocs_host` information, as shown in the following figure.



```
$ sudo docker inspect dad1ad2379ab | grep ocs
  "_AUTOCONF_ocs__host=a1d2af7272c64107.m.cnhzaligrpzmfpub001.ocs.aliyuncs.com",
  "_AUTOCONF_ocs__name=TODO",
  "_AUTOCONF_ocs__password=wangcuiFY102300",
  "_AUTOCONF_ocs__port=11211",
  "_AUTOCONF_ocs__username=a1d2af7272c64107",
```

`host` is a list of servers on which OCS (caching service) is deployed. `port` indicates the port number.

Machine Learning Platform for AI in Apsara Stack typically uses the built-in memcached service as the dependent caching service. If port 11211 is inaccessible, log on to the server and run the following command to restart the memcached service:

```
docker restart containerid
```

3.8.2.1.2.3 Experiment failures

We recommend that you run a Machine Learning Platform for AI experiment in Google Chrome version 66 or later. Google Chrome is the only supported browser.

- Components cannot be dragged and dropped.

Clear cookies and caches, and then retry. Check the version of Chrome. If the problem persists, it is due to a service failure. Log on to the container to view the log.

- An error message is displayed while an algorithm is running.

If an error message is displayed, the task has been submitted to MaxCompute. Check the parameters and source data against the user guide and algorithm descriptions to locate the error.

3.8.2.1.2.4 Other failures

If a problem persists after you have checked all items by referring to [pai.xx.xx access failures](#), troubleshoot the underlying dependency services, including MaxCompute and DataWorks (tenants and metadata).

- **MaxCompute:** Make sure that MaxCompute can pass the `pai_console` test.
- **DataWorks:** Make sure the configured domain name is accessible, and verify the application log.

If no errors are found, restart the service.

3.8.2.2 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.

3.8.2.3 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 `runctl` 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 `r al` 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.

3.9 Elasticsearch

3.9.1 Log on to the Apsara Stack Elasticsearch O&M center

Procedure

1. In the left-side navigation pane, choose **Products > Product List > Apsara Bigdata Manager**.

2. In the Apsara Bigdata Manager center, choose **Monitor > Dashboard**, and click **Elasticsearch** in the **Summary** section to log on to the Apsara Stack Elasticsearch O&M center.

3.9.2 Apsara Stack Elasticsearch O&M center

3.9.2.1 O&M overview

This topic describes the features of Elasticsearch O&M and how to access the Elasticsearch O&M page.

Modules

Elasticsearch O&M includes **business O&M**, **service O&M**, **cluster O&M**, and **host O&M**. The following table describes them in detail.

Module	Feature	Description
Business O&M	Cluster Configuration	Allows you to view and modify the cluster configuration files of the worker and kibana nodes for Elasticsearch.
	System Configuration	Allows you to view and modify the system configuration files for Elasticsearch.
Service O&M	Overview	Displays all Elasticsearch services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission , TCP connection, and root disk usage for each service.
	Hosts	Displays all hosts where each Elasticsearch service is run so that you can understand the service deployment on hosts.
Cluster O&M	Overview	Displays the overall running and health check information about a cluster. On this page, you can view the host status, service status, health check result, and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, heap memory usage, TCP connection, and root disk usage.

Module	Feature	Description
	Health Status	Displays all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.
Host O&M	Overview	Displays the overall running and health check information about a host. On this page, you can view the root disk usage, total usage, 1-minute load, 5-minute load, 15-minute load, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage.
	Charts	Displays the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.
	Health Status	Displays the check results for a host. The check results are divided into the Critical, Warning, Exception, and OK types.
	Services	Displays information about service instances and service instance roles of a host.

Entry

1. Log on to the ABM console.
2. Click  in the upper-left corner, and then click Elasticsearch.
3. On the Elasticsearch page that appears, click O&M in the upper-right corner. The Business page appears.

The O&M page includes four modules, namely, Business, Services, Clusters, and Hosts.

3.9.2.2 Business O&M

3.9.2.2.1 Cluster configuration

This topic describes how to view and modify the cluster configuration files of the worker and kibana nodes for Elasticsearch in Apsara Bigdata Manager (ABM).

Entry

1. At the top of the O&M page, click the Business tab.
2. On the Business page that appears, click Cluster Configuration in the left-side navigation pane.
3. In the worker or kibana list, click a cluster configuration file that you want to view. The details of the file appear on the right.

Modify a cluster configuration file

1. Click a cluster configuration file to be modified and click Edit to modify the configuration file.
2. Click Save.
3. Click Preview.
 - a. In the Preview dialog box that appears, you can compare the differences before and after the file modification.
 - b. If the modification is correct, click OK.
4. Click Submit at the bottom of the page. The modification is completed.

If you want to undo the modification, click Undo.

Upload a plug-in



Notice:

The custom plug-in may affect the stability of the cluster. Make sure that the custom plug-in is reliable and secure to use. The plug-in is not automatically updated with Elasticsearch. To update the plug-in, you must manually upload a new version of the plug-in.

1. Select a cluster to which you want to upload a plug-in from the drop-down list. Click Upload Plug-in.
2. In the Upload Plug-in dialog box that appears, click Click here to select files for upload to upload one or more files.

To delete a file that no longer needs to be uploaded, click x next to the file.

3.  **Notice:**

The custom plug-in may affect the stability of the cluster. Make sure that the custom plug-in is reliable and secure to use. The plug-in is not automatically updated with Elasticsearch. To update the plug-in, you must manually upload a new version of the plug-in.

Select the check box in the dialog box.

4. Click OK.

3.9.2.2.2 System configuration

This topic describes how to view and modify the system configuration files for Elasticsearch in Apsara Bigdata Manager (ABM).

Entry

1. At the top of the O&M page, click the Business tab.
2. On the Business page that appears, click System Configuration in the left-side navigation pane.
3. Click a configuration file that you want to view. The details of the file appear on the right.

Modify a system configuration file

1. Click a system configuration file to be modified and click Edit to modify the configuration file.
2. Click Save.
3. Click Preview.
 - a. In the Preview dialog box that appears, you can compare the differences before and after the file modification.
 - b. If the modification is correct, click OK.
4. Click Submit at the bottom of the page. The modification is completed.

If you want to undo the modification, click Undo.

3.9.2.3 Service O&M

3.9.2.3.1 Service overview

The service overview page lists all Elasticsearch services in a cluster. You can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for each service.

Entry

1. At the top of the O&M page, click the Services tab.
2. On the Services page that appears, select a service in the left-side navigation pane. Click the Overview tab.

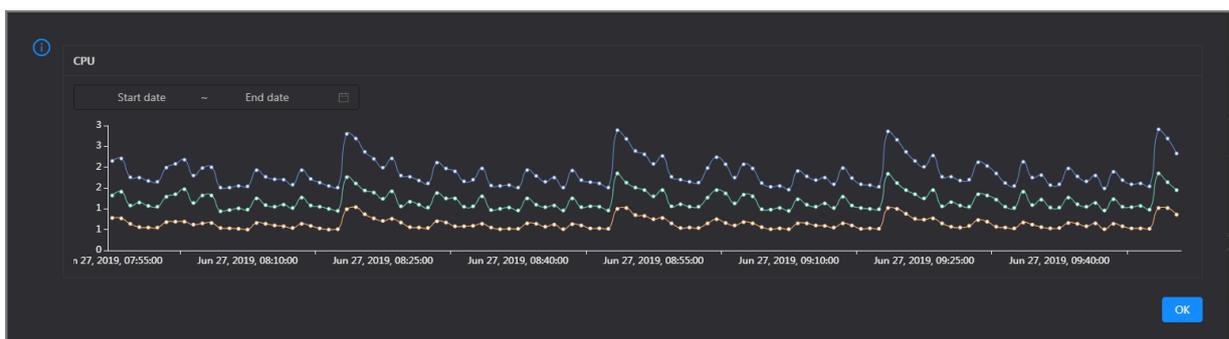
On the Overview page that appears, you can view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, TCP connection, and root disk usage for the selected service.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

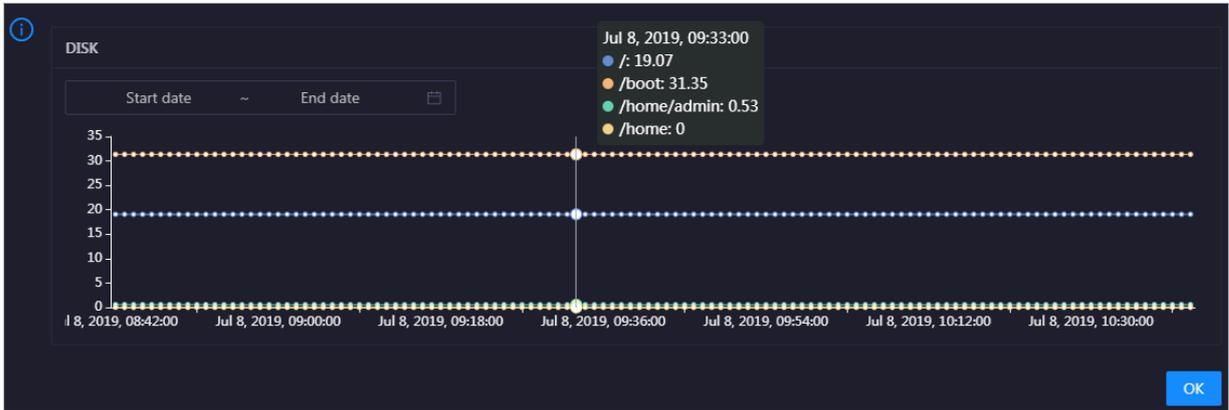
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the service in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

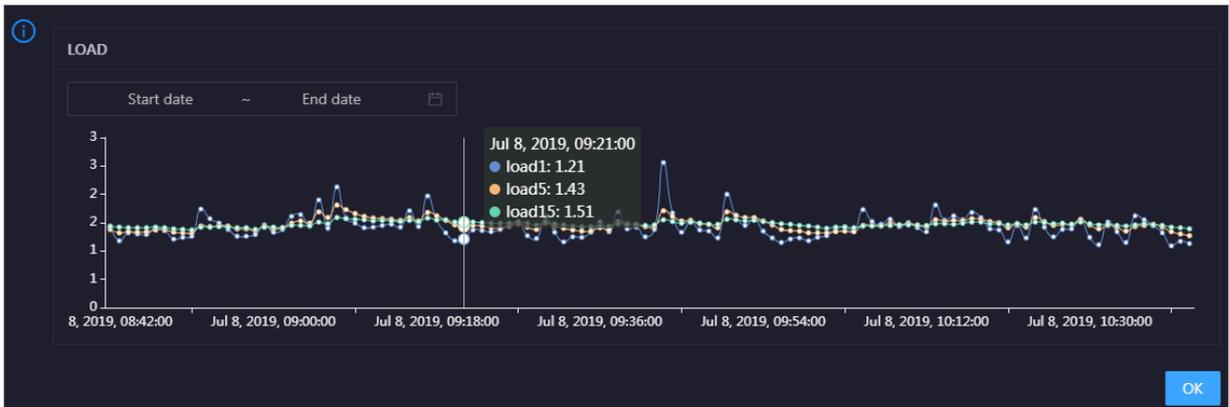


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage space usage of the service in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

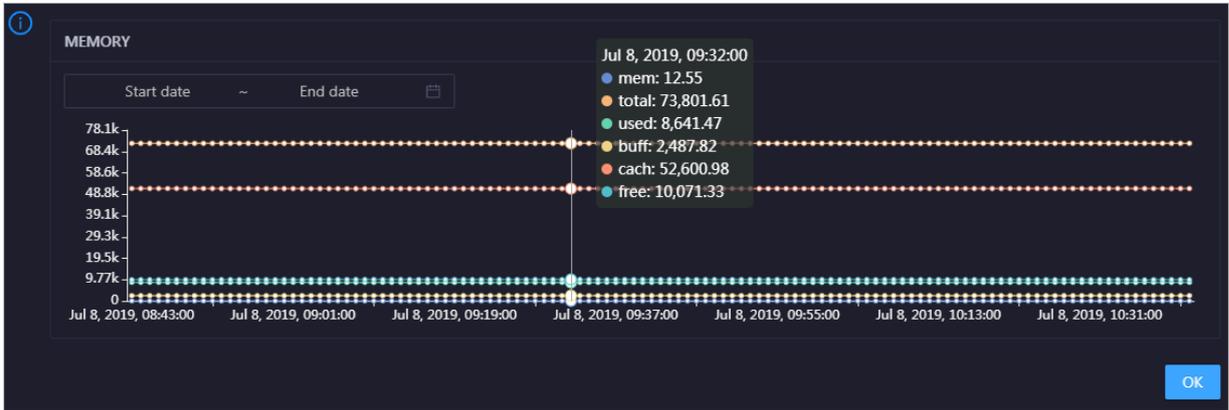


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the selected service in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the selected service over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

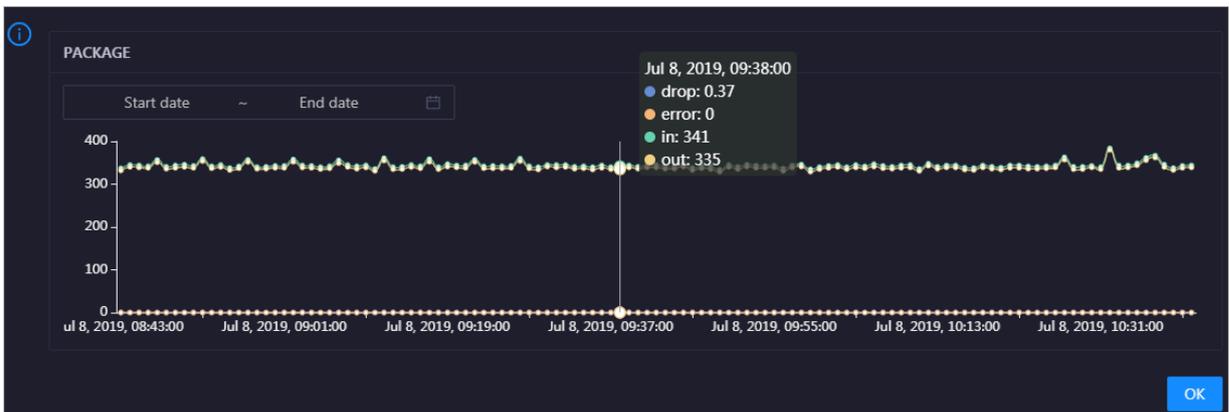


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the selected service in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the selected service over time in different colors. These trend lines reflect the data transmission status of the service.

Click  in the upper-right corner of the chart to zoom in it.



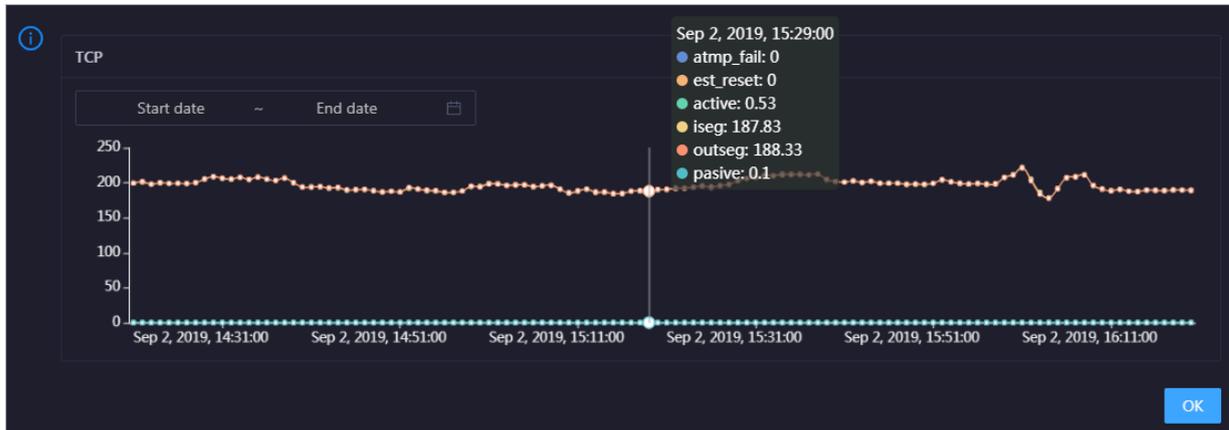
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the selected service in the specified period.

TCP

This chart displays the trend lines of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP

packets (outseg) for the selected service over time in different colors. These trend lines reflect the TCP connection status of the service.

Click  in the upper-right corner of the chart to zoom in it.

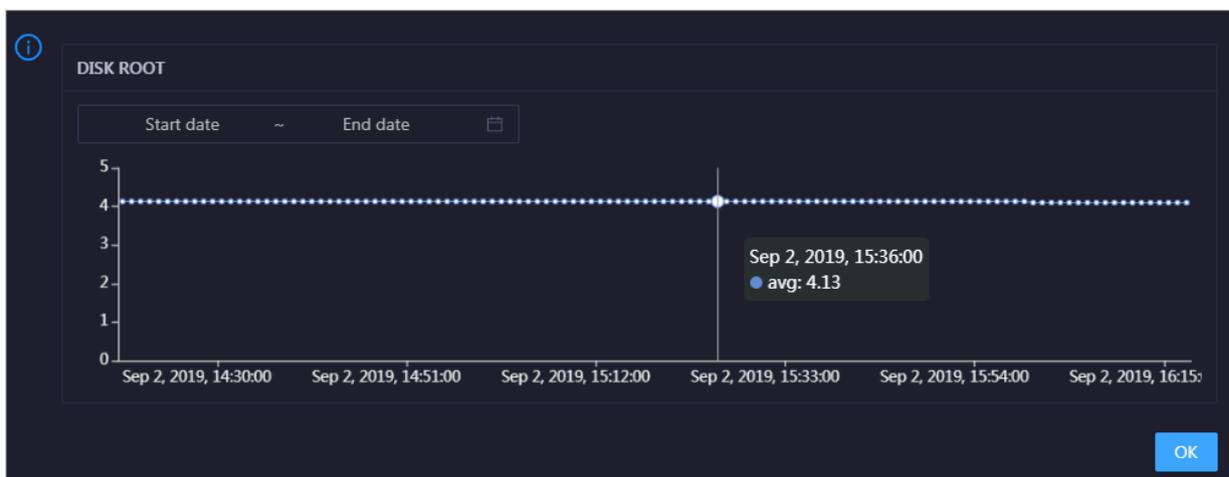


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the selected service in the specified period.

DISK ROOT

This chart displays the trend line of the average root disk usage (avg) for the selected service over time.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the selected service in the specified period.

3.9.2.3.2 Service hosts

This topic describes how to view all hosts where each Elasticsearch service is run.

On the Server page, you can view the hosts where the selected service is run.

1. At the top of the O&M page, click the Services tab.
2. On the Services page that appears, select a service in the left-side navigation pane.
3. Click the Server tab. The Server page for the service appears.

On the Server page, you can view the hosts where the selected service is run.

3.9.2.4 Cluster O&M

3.9.2.4.1 Cluster overview

The cluster overview page displays the overall running and health check information about a cluster. On this page, you can view the host status, service status, health check result, and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, heap memory usage, TCP connection, and root disk usage.

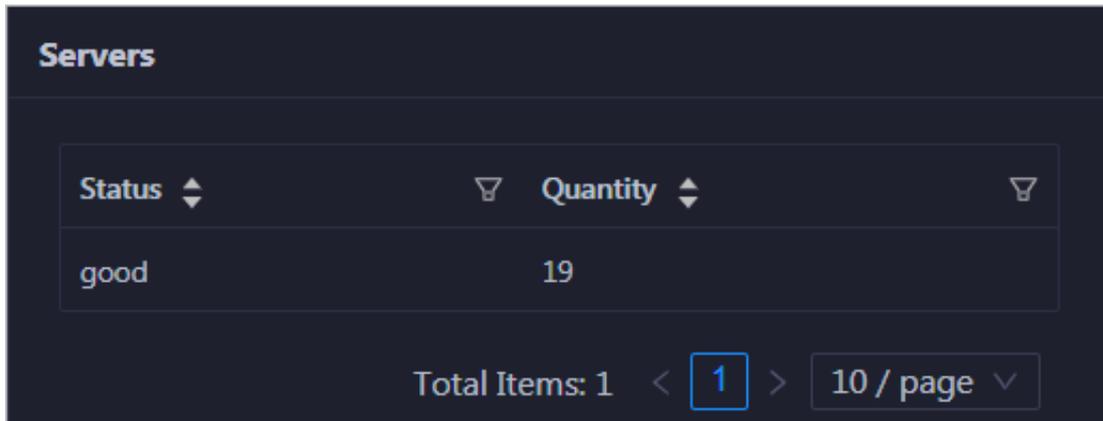
Entry

1. At the top of the O&M page, click the Clusters tab.
2. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Overview tab. The Overview page for the cluster appears.

On the Overview page, you can view the host status, service status, health check result, and health check history of the cluster. You can also view the trend charts of CPU usage, disk usage, memory usage, load, packet transmission, heap memory usage, TCP connection, and root disk usage. To view information about a cluster, select a region in the left-side navigation pane, and then select the cluster in the region.

Servers

This section displays all host statuses and the number of hosts in each status. The host statuses include good and bad.



Services

This section displays all services deployed in the cluster and the respective number of available and unavailable services.

Health Check

This section displays the number of checkers deployed for the cluster and the respective number of Critical, Warning, and Exception alerts.



Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

Health Check History

This section displays a record of the health checks performed on the cluster.

Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see [Cluster health](#).

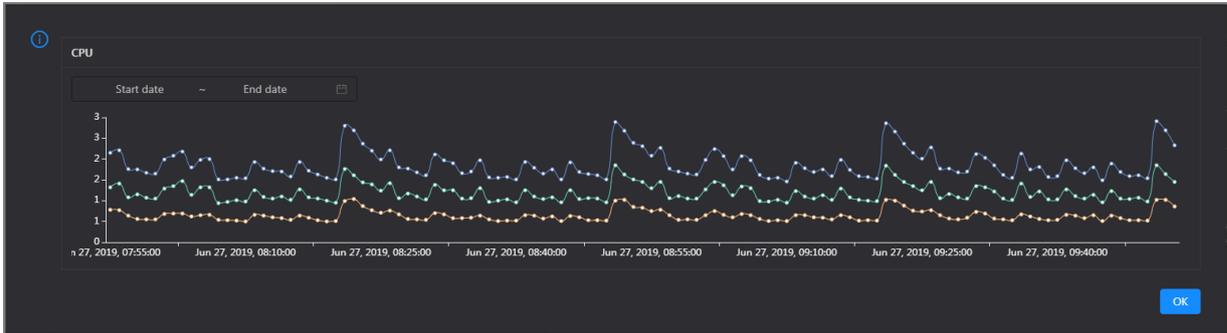
You can click the event content of a check to view the exception items.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

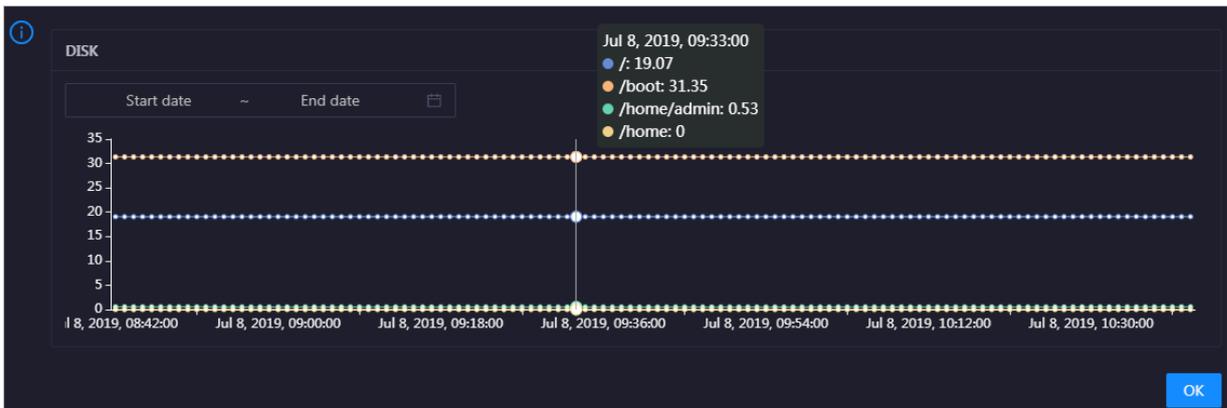
You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the cluster in the specified period.



DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

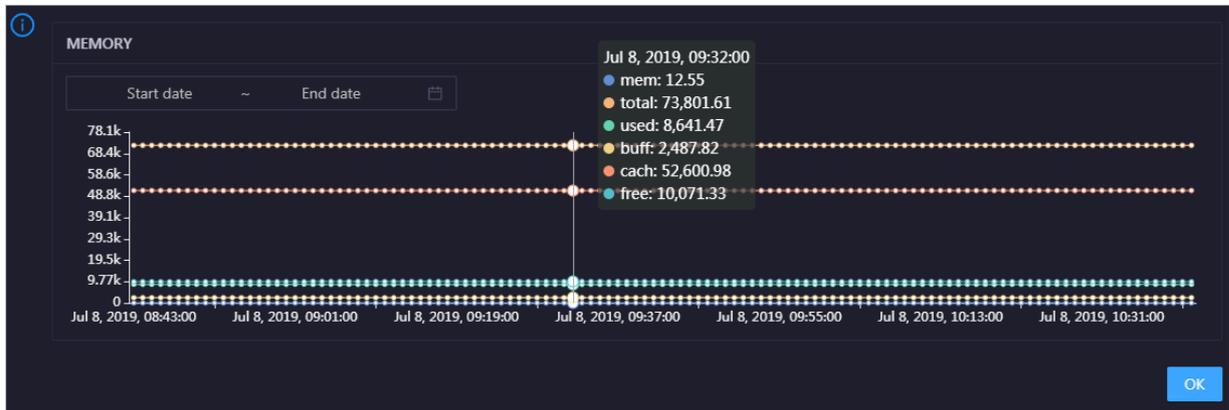


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the cluster in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (mem), total memory size (total), used memory size (used), size of memory used by kernel buffers (buff), size of memory used by the page cache (cach), and available memory size (free) for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

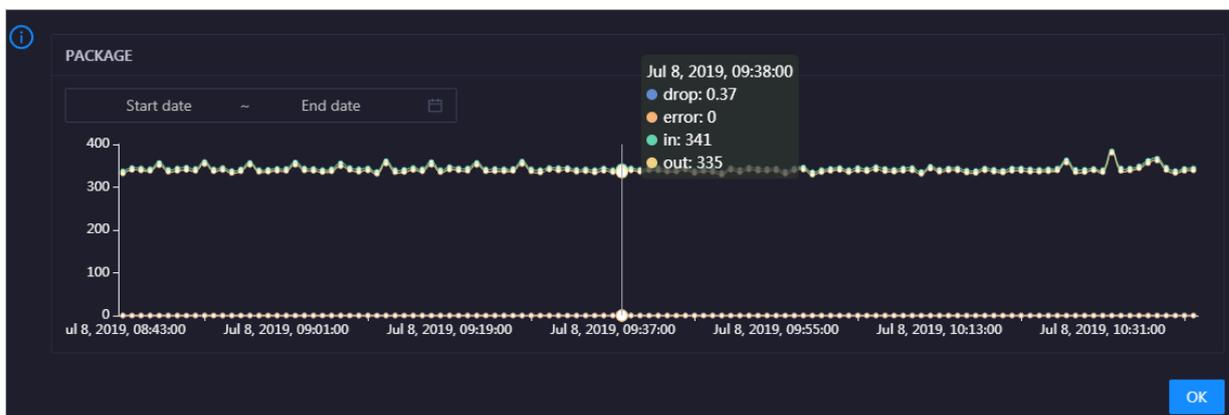


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the cluster in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the cluster over time in different colors. These trend lines reflect the data transmission status of the cluster.

Click  in the upper-right corner of the chart to zoom in it.

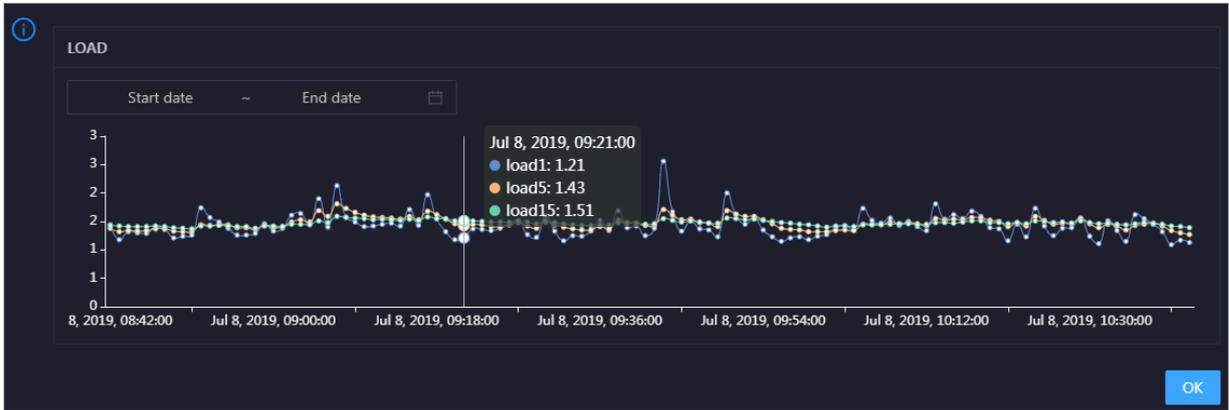


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the cluster in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the cluster over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the cluster in the specified period.

HEAP

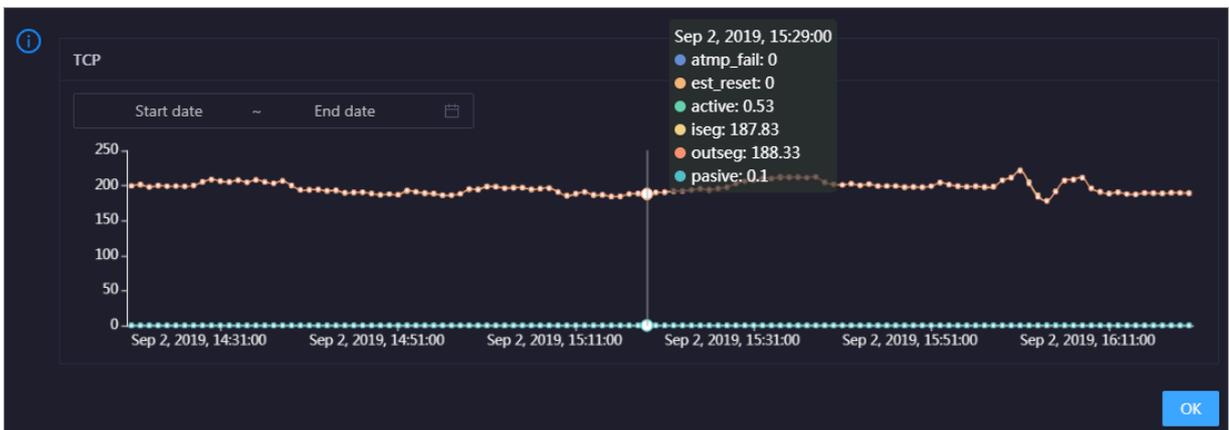
This chart displays the trend lines of the heap memory usage over time in different colors.

Click  in the upper-right corner of the chart to zoom in the chart.

TCP

This section displays the trend line of the number of failed TCP connection attempts (atmp_fail), that of the times of resetting TCP connections in the ESTABLISHED state (est_reset), that of active TCP connections (active), that of passive TCP connections (pasive), that of received TCP packets (iseg), and that of sent TCP packets (outseg) for the cluster over time in different colors. These trend lines reflect the TCP connection status of the cluster.

Click  in the upper-right corner of the chart to zoom in the chart.

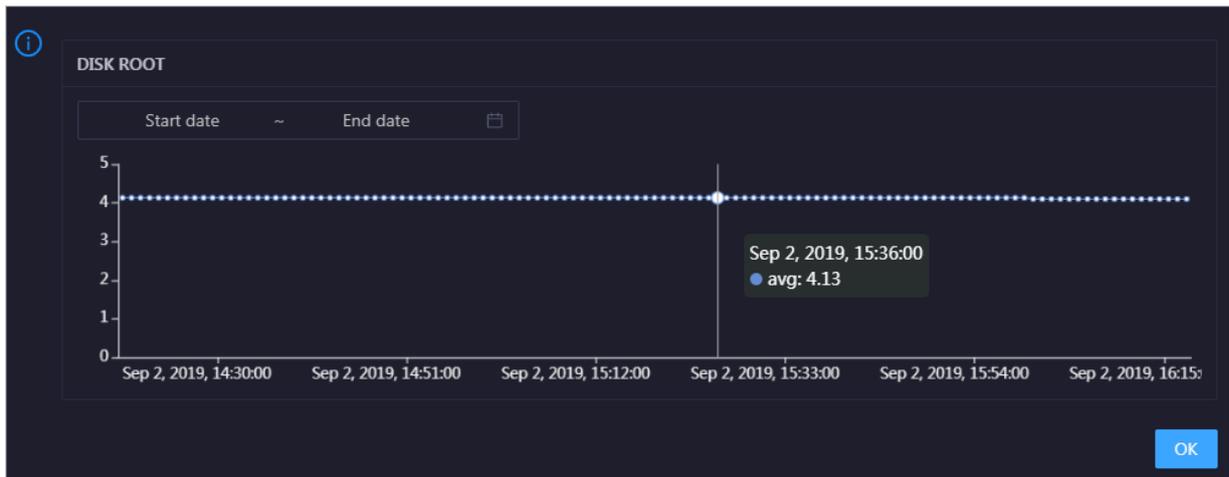


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the TCP connection status of the cluster in the specified period.

DISK ROOT

This section displays the trend line of the average usage of the root disk (/) for the cluster over time.

Click  in the upper-right corner of the chart to zoom in the chart.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the average root disk usage of the cluster in the specified period.

3.9.2.4.2 Cluster health

On the cluster health status page, you can view all checkers of a cluster, including the checker details, check results for the hosts in the cluster, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

Entry

At the top of the O&M page, click the Clusters tab. On the Clusters page, select a cluster in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the cluster appears.

Checker	Source	Critical	Warning	Exception	Actions
+ bcc_host_live_check	tcheck	3	0	0	Details
+ elasticsearch_check_health_shuttle	tcheck	2	0	0	Details
+ bcc_check_ntp	tcheck	0	0	0	Details
+ bcc_tsar_tcp_checker	tcheck	0	0	0	Details
+ bcc_kernel_thread_count_checker	tcheck	0	0	0	Details
+ bcc_network_tcp_connections_checker	tcheck	0	0	0	Details
+ bcc_disk_usage_checker	tcheck	0	0	0	Details
+ bcc_process_thread_count_checker	tcheck	0	0	0	Details
+ bcc_check_load_high	tcheck	0	0	0	Details

On the Health Status page, you can view all checkers of the cluster and the check results for the hosts in the cluster. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.

Name: bcc_tsar_tcp_checker **Source:** tcheck

Alias: TCP Retransmission Check **Application:** bcc

Type: system **Scheduling:** Enable

Data Collection: Enable

Default Execution Interval: 0 0/5 * * * ?

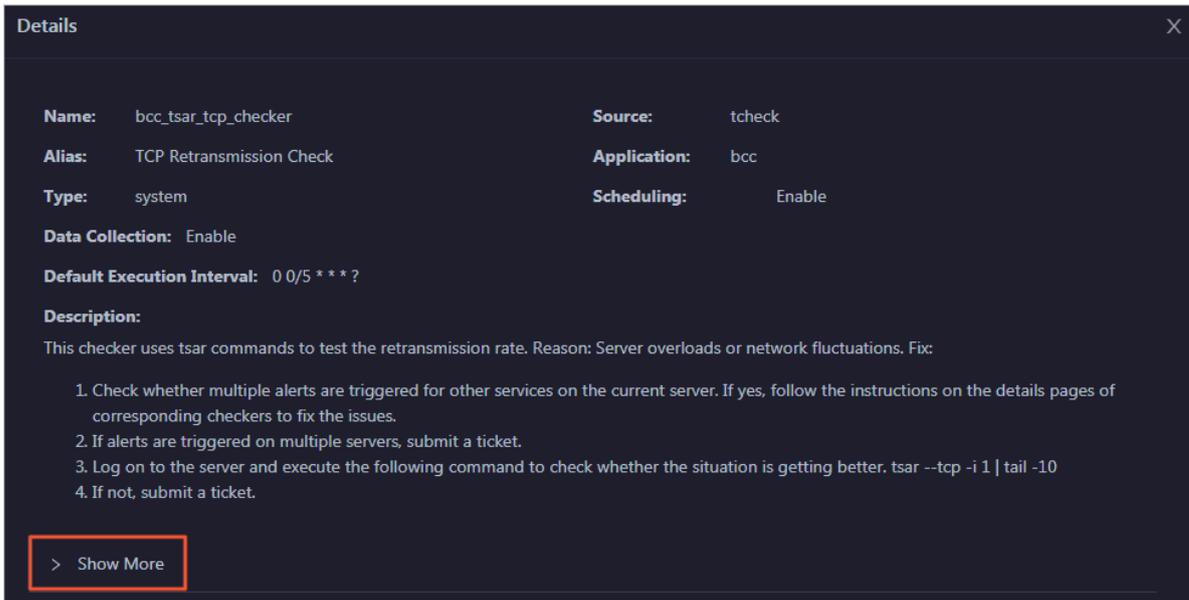
Description:
 This checker uses tsar commands to test the retransmission rate. Reason: Server overloads or network fluctuations. Fix:

1. Check whether multiple alerts are triggered for other services on the current server. If yes, follow the instructions on the details pages of corresponding checkers to fix the issues.
2. If alerts are triggered on multiple servers, submit a ticket.
3. Log on to the server and execute the following command to check whether the situation is getting better. `tsar --tcp -i 1 | tail -10`
4. If not, submit a ticket.

> Show More

The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

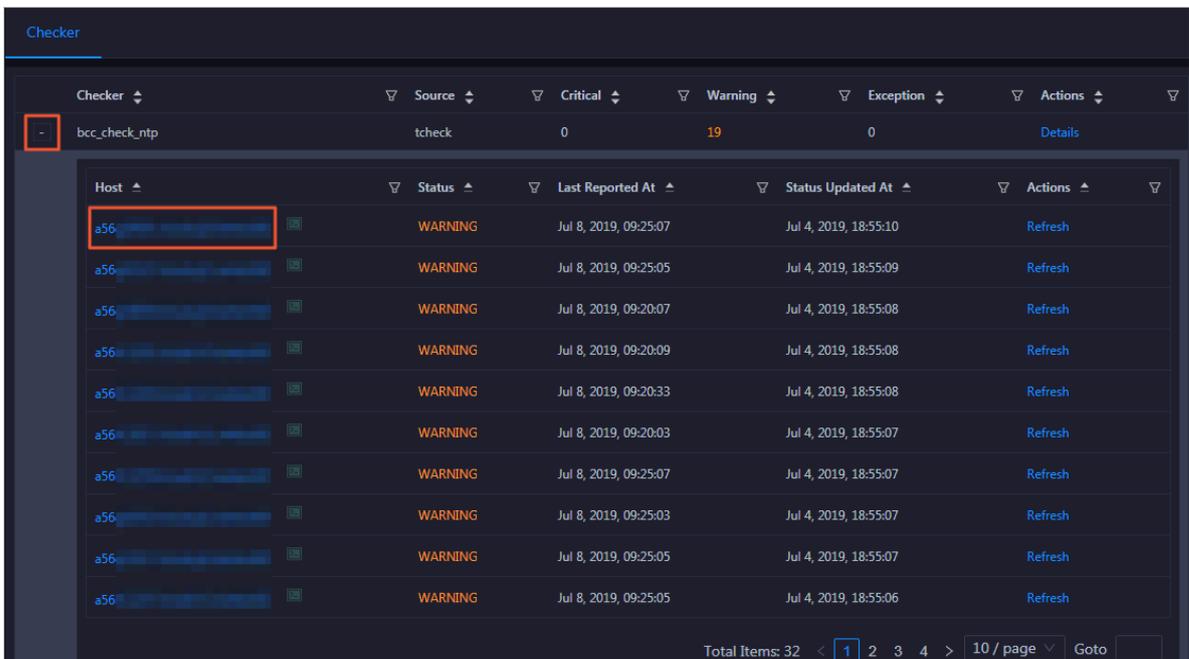


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

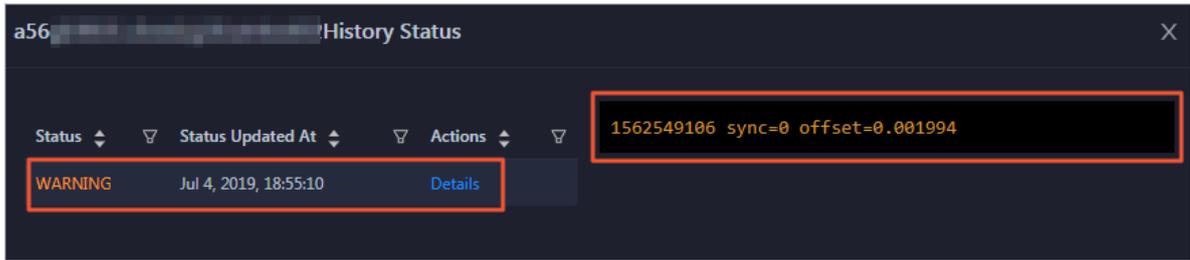
View hosts with alerts and the alert causes

You can view the check history and check results of a checker on a host.

1. On the Health Status page, click + to expand a checker with alerts. You can view all hosts where the checker is run.

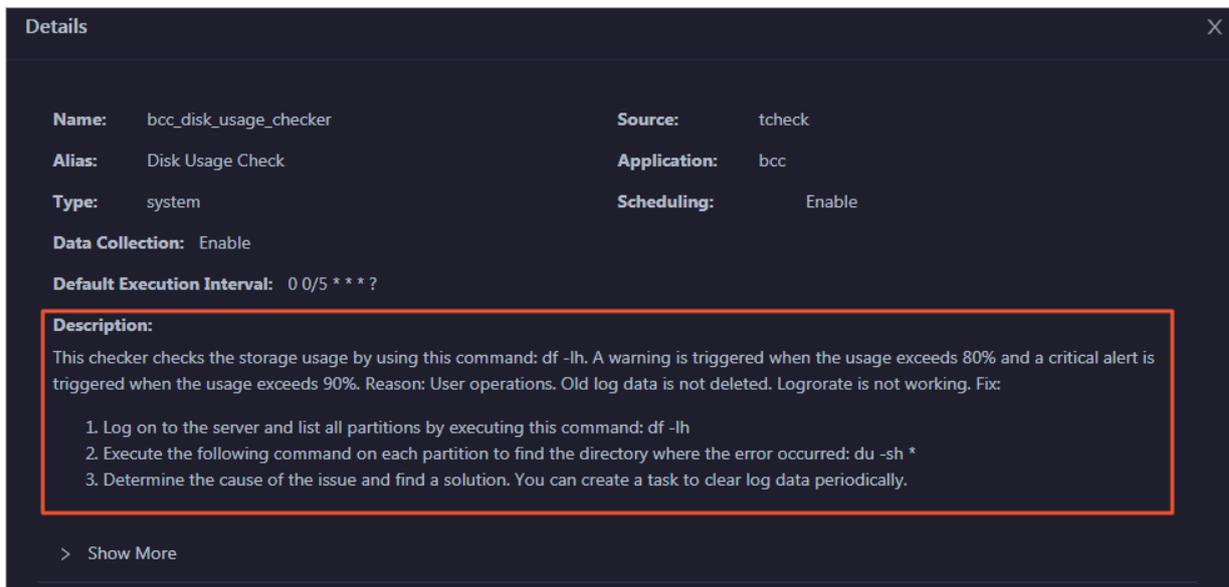


2. Click a hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

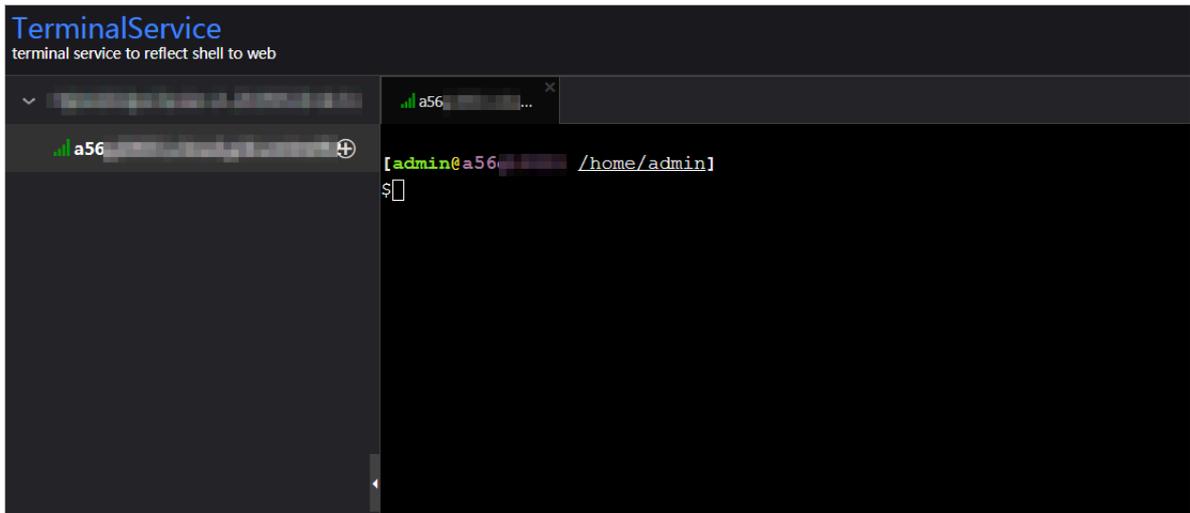
- On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



Log on to a host

- To log on to a host to clear alerts or perform other operations, follow these steps:

3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.

Checker	Source	Critical	Warning	Exception	Actions
bcc_check_ntp	tcheck	0	19	0	Details

Host	Status	Last Reported At	Status Updated At	Actions
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:10	Refresh
a56	WARNING	Jul 8, 2019, 09:25:05	Jul 4, 2019, 18:55:09	Refresh
a56	WARNING	Jul 8, 2019, 09:20:07	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:09	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:33	Jul 4, 2019, 18:55:08	Refresh
a56	WARNING	Jul 8, 2019, 09:20:03	Jul 4, 2019, 18:55:07	Refresh
a56	WARNING	Jul 8, 2019, 09:25:07	Jul 4, 2019, 18:55:07	Refresh

3.9.2.5 Host O&M

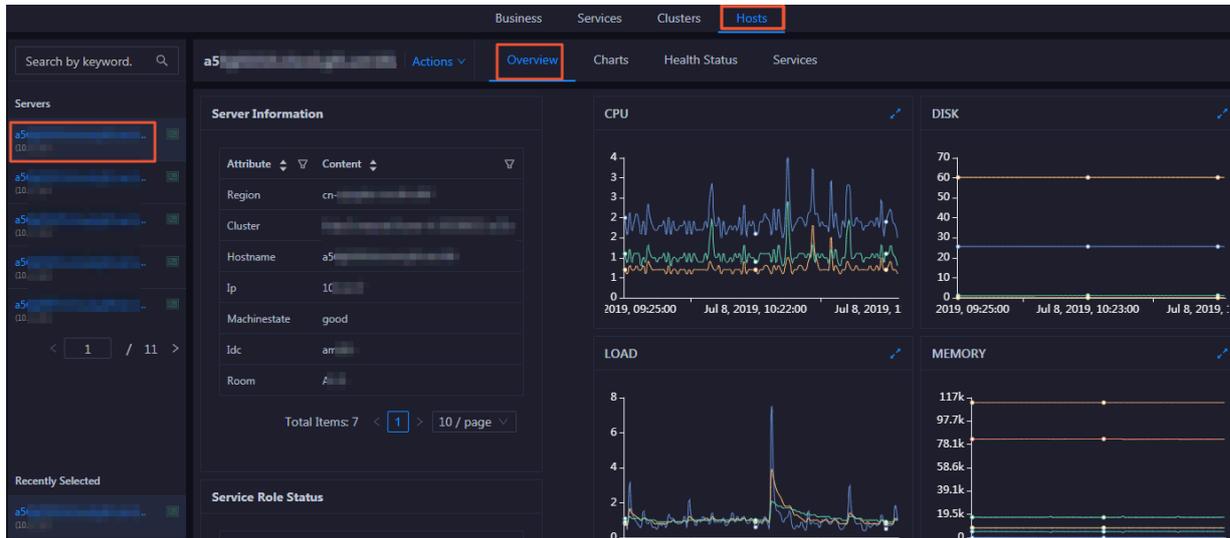
3.9.2.5.1 Host overview

The host overview page displays the overall running information about a host in an Elasticsearch cluster. On this page, you can view the information, service role status, health check result, and health check history of the host. You can also

view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

Entry

On the Hosts page, select a host in the left-side navigation pane, and then click the Overview tab. The Overview page for the host appears.



On the Overview page, you can view the server information, service roles, health check result, and health check history of the host. You can also view the trend charts of CPU usage, disk usage, memory usage, load, and packet transmission.

Server Information

This section displays the information about the host, including the region, cluster, name, IP address, status, IDC, and server room of the host.

Server Information

Attribute	Content
Region	cn-██████████
Cluster	██████████
Hostname	a56-██████████
Ip	10.██████
Machinestate	good
Idc	am-██████
Room	A-██████

Total Items: 7 < 1 > 10 / page

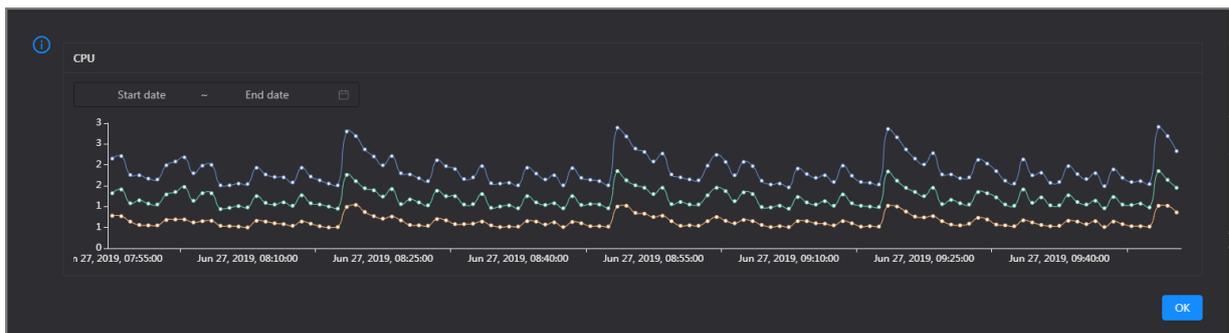
Service Role Status

This section displays the information about the services deployed on the host, including the roles, statuses, and number of services.

CPU

This chart displays the trend lines of the total CPU usage (cpu), CPU usage for executing code in kernel space (sys), and CPU usage for executing code in user space (user) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

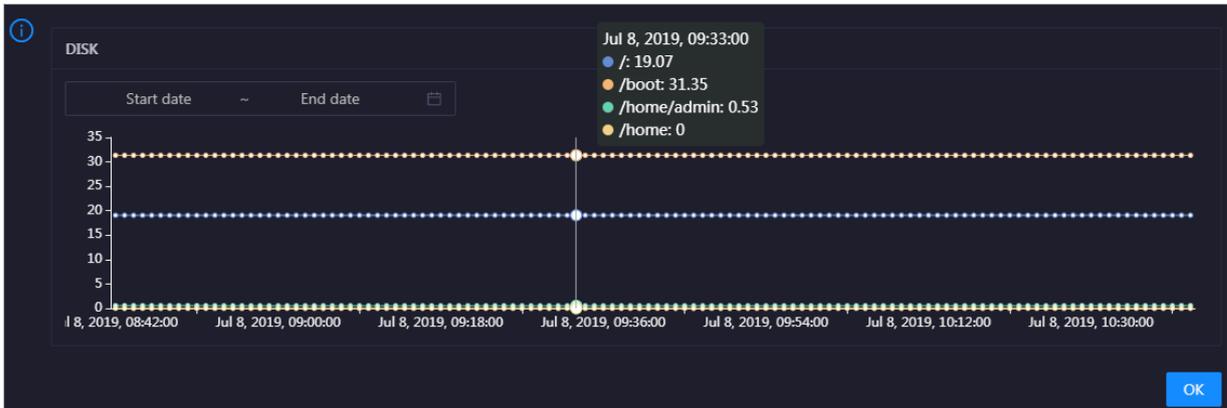


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the CPU usage of the host in the specified period.

DISK

This chart displays the trend lines of the storage space usage on the `/`, `/boot`, `/home/admin`, and `/home` directories for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

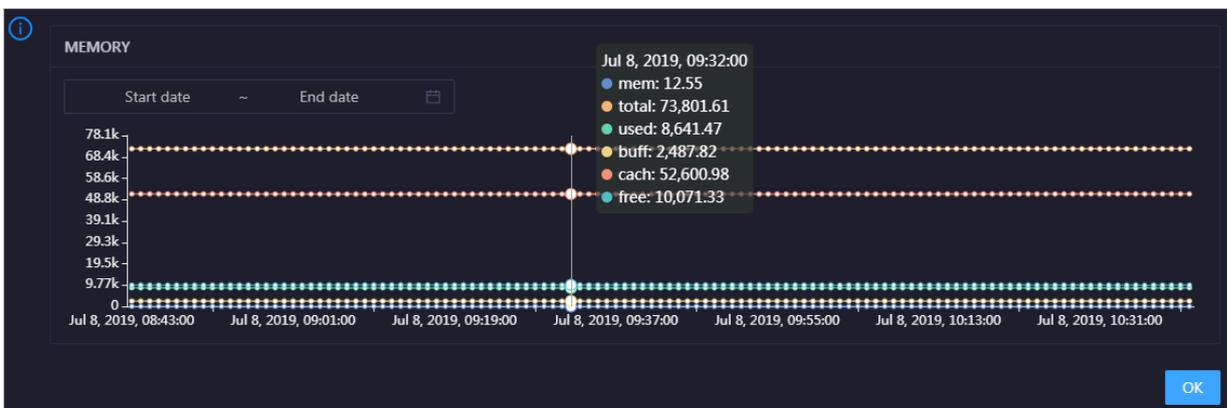


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the storage usage of the host in the specified period.

MEMORY

This chart displays the trend lines of the memory usage (`mem`), total memory size (`total`), used memory size (`used`), size of memory used by kernel buffers (`buff`), size of memory used by the page cache (`cach`), and available memory size (`free`) for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

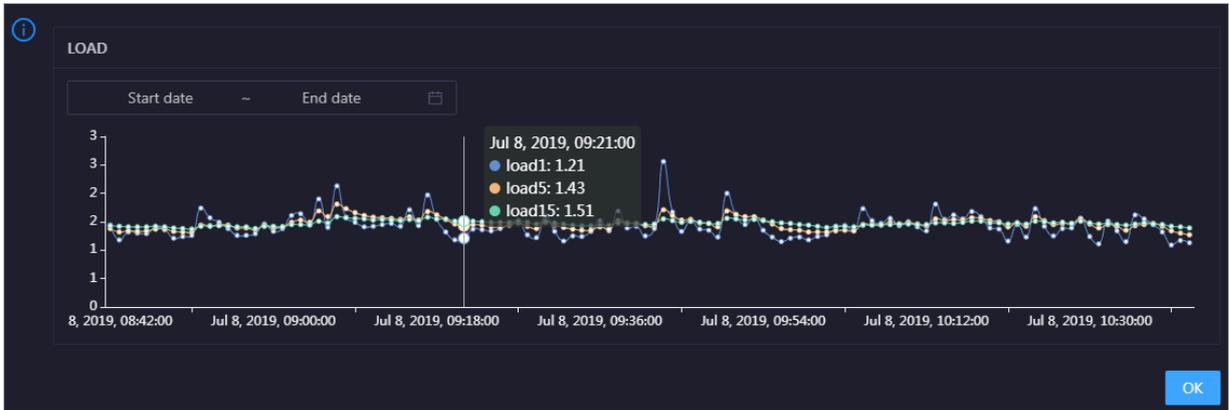


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the memory usage of the host in the specified period.

LOAD

This chart displays the trend lines of the 1-minute, 5-minute, and 15-minute load averages for the host over time in different colors.

Click  in the upper-right corner of the chart to zoom in it.

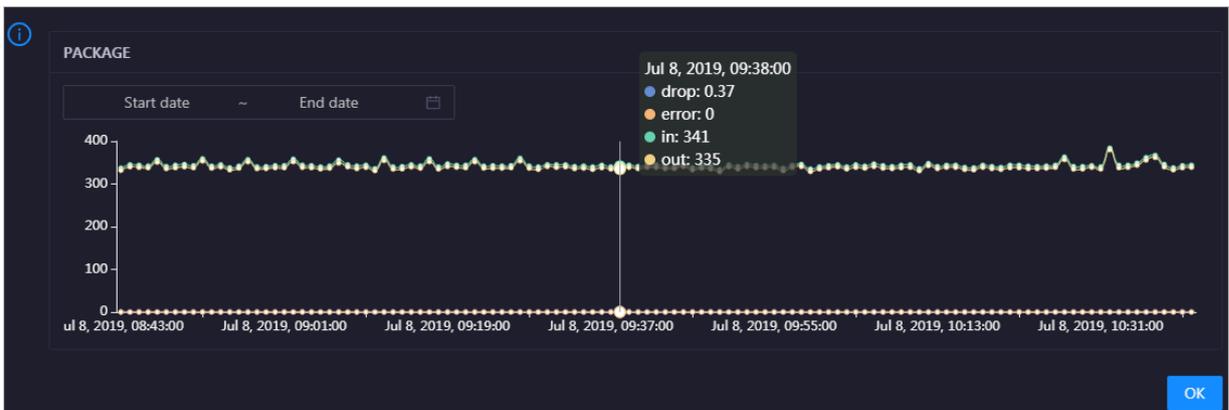


You can specify the start time and end time in the upper-left corner of the enlarged chart to view the 1-minute, 5-minute, and 15-minute load averages of the host in the specified period.

PACKAGE

This chart displays the trend lines of the number of dropped packets (drop), that of error packets (error), that of received packets (in), and that of sent packets (out) for the host over time in different colors. These trend lines reflect the data transmission status of the host.

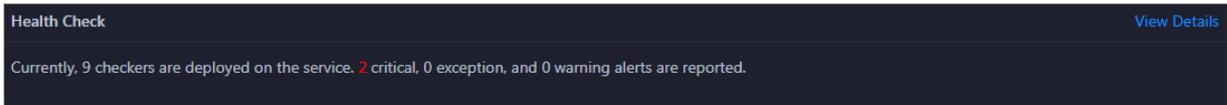
Click  in the upper-right corner of the chart to zoom in it.



You can specify the start time and end time in the upper-left corner of the enlarged chart to view the data transmission status of the host in the specified period.

Health Check

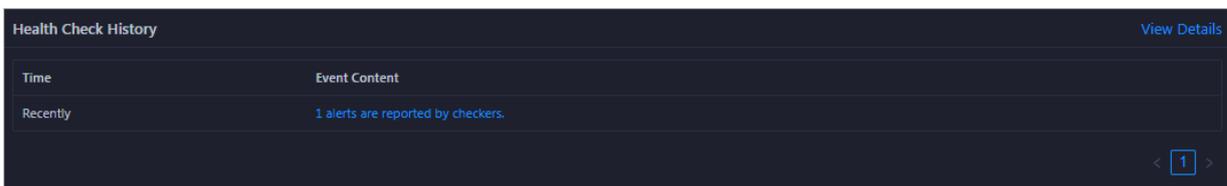
This section displays the number of checkers deployed for the host and the respective number of Critical, Warning, and Exception alerts.



Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

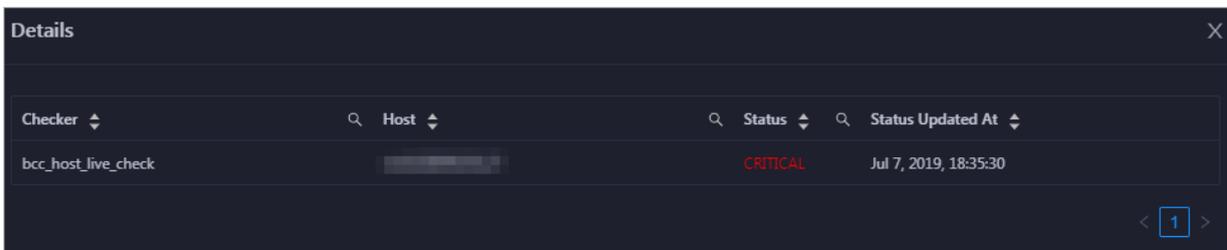
Health Check History

This section displays a record of the health checks performed on the host.



Click View Details to go to the Health Status page. On this page, you can view the health check details. For more information, see [Host health](#).

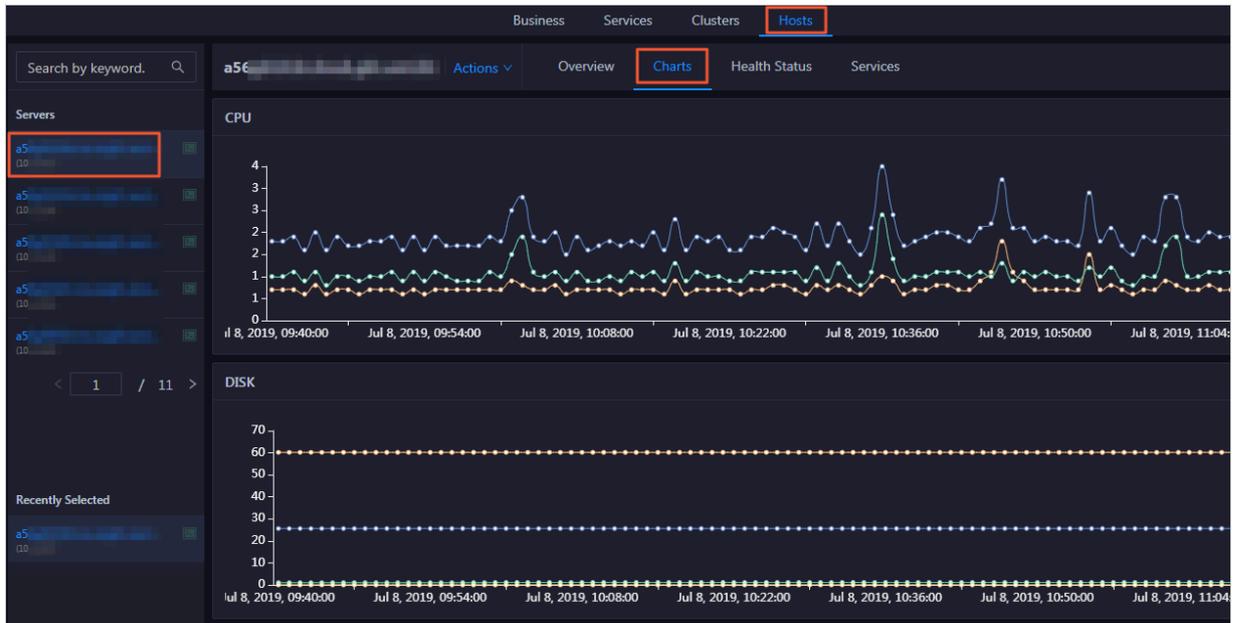
You can click the event content of a check to view the exception items.



3.9.2.5.2 Host charts

On the host chart page, you can view the enlarged trend charts of CPU usage, memory usage, storage usage, load, and packet transmission.

On the Hosts page, select a host in the left-side navigation pane, and then click the Charts tab. The Charts page for the host appears.



The Charts page displays trend charts of CPU usage, disk usage, memory usage, load, and packet transmission for the host. For more information, see [Host overview](#).

3.9.2.5.3 Host health

On the host health status page, you can view the checkers of the selected host, including the checker details, check results, check history, and schemes to clear alerts (if any). In addition, you can log on to a host and perform manual checks on the host.

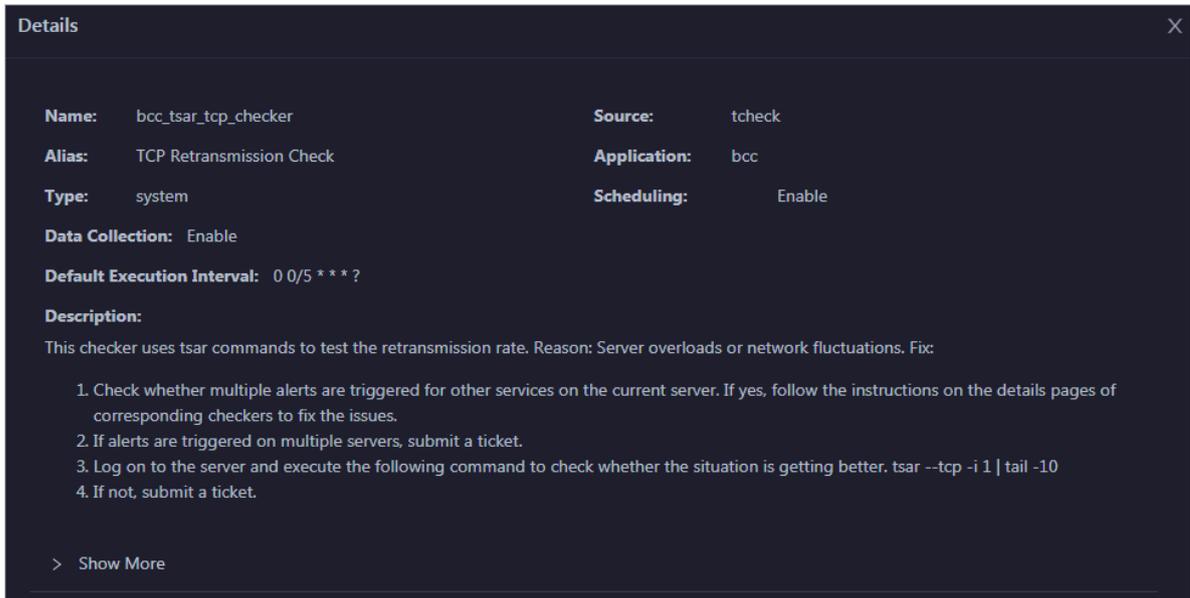
Entry

At the top of the O&M page, click the Hosts tab. On the page that appears, select a host in the left-side navigation pane, and then click the Health Status tab. The Health Status page for the host appears.

On the Health Status page, you can view all checkers and the check results for the host. The check results are divided into the Critical, Warning, Exception, and OK types. They are displayed in different colors. Among them, the Critical, Warning, and Exception results are alerts. You need to pay attention to them, especially the Critical and Warning results.

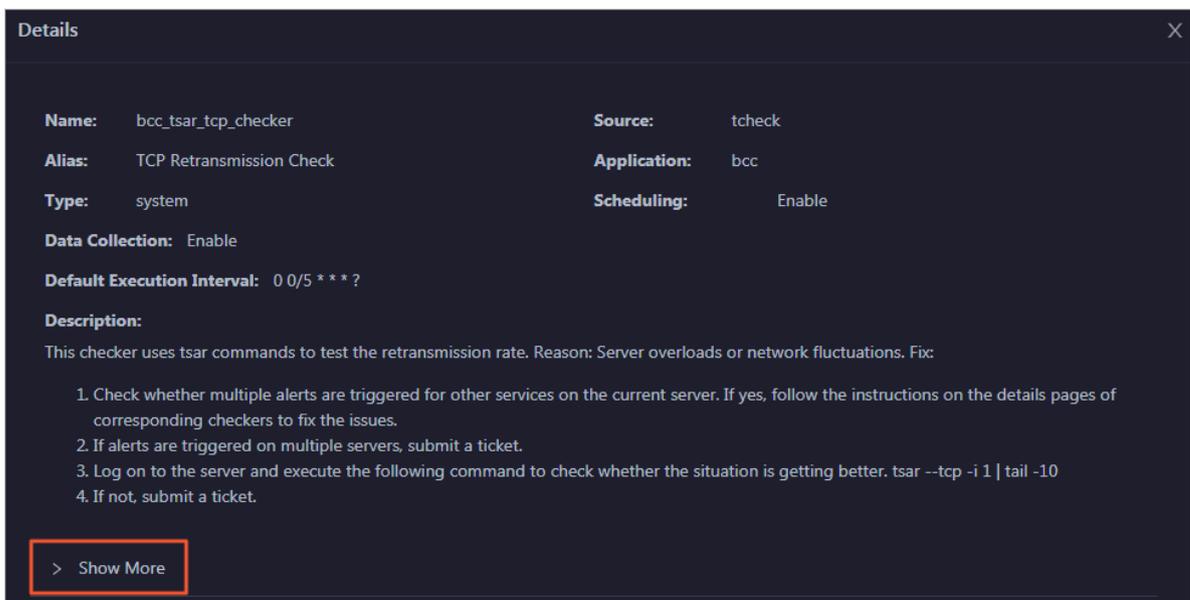
View checker details

1. On the Health Status page, click Details in the Actions column of a checker. In the dialog box that appears, view the checker details.



The checker details include the name, source, alias, application, type, default execution interval, and description of the checker, whether scheduling is enabled, and whether data collection is enabled. The schemes to clear alerts are provided in the description.

2. Click Show More at the bottom to view more information about the checker.

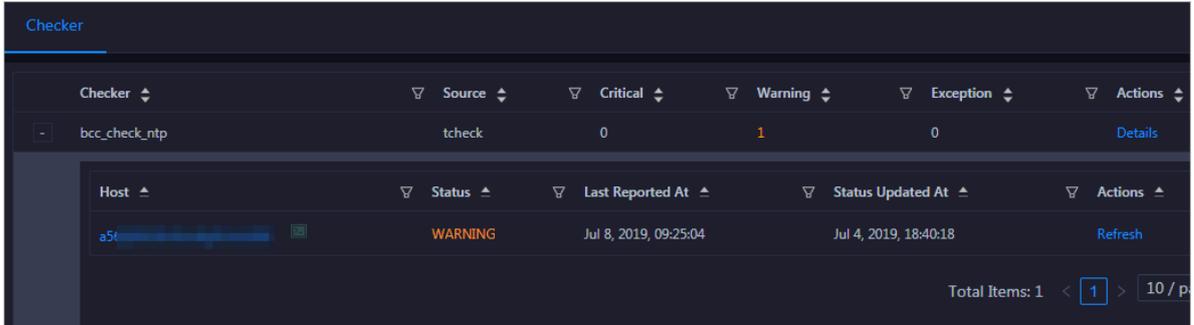


You can view information about the execution script, execution target, default threshold, and mount point for data collection.

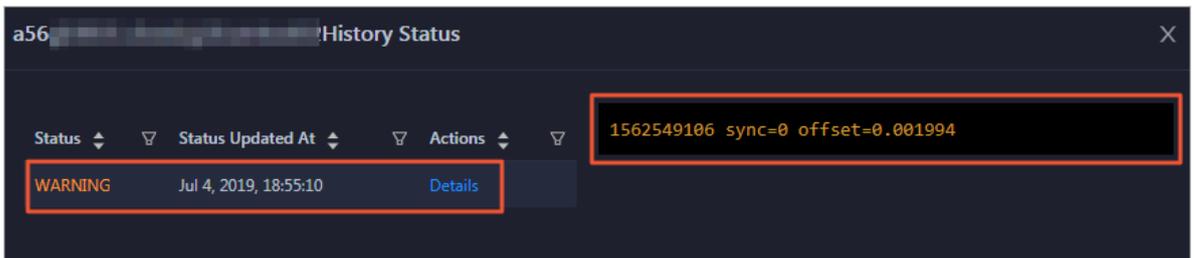
View alert causes

You can view the check history and check results of a checker.

1. On the Health Status page, click + to expand a checker with alerts.

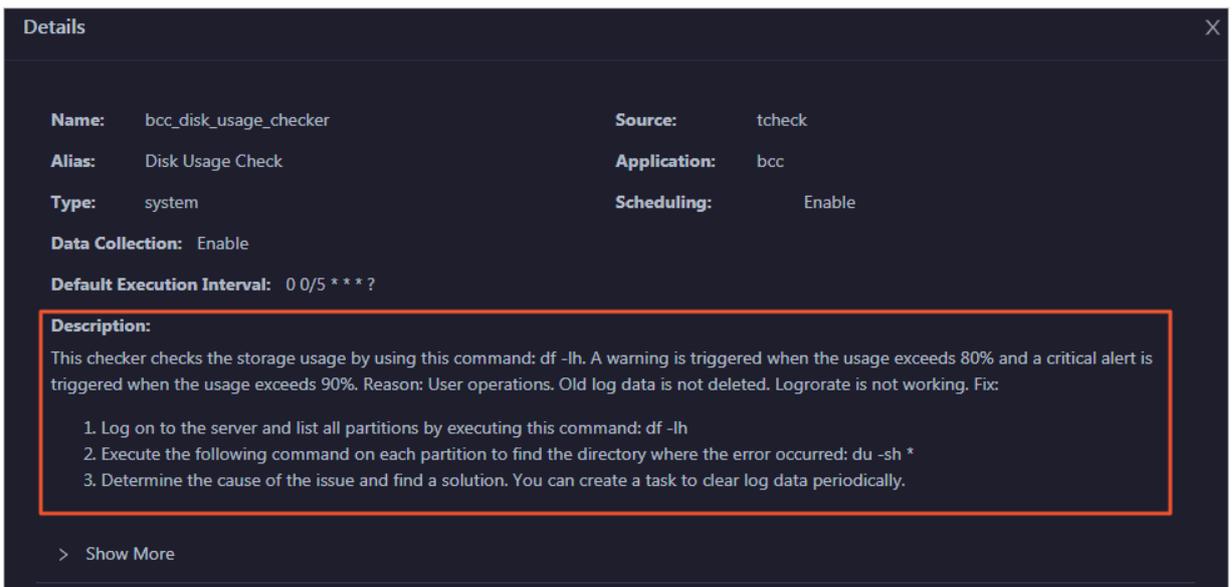


2. Click the hostname. In the dialog box that appears, click Details in the Actions column of a check result to view the alert causes.



Clear alerts

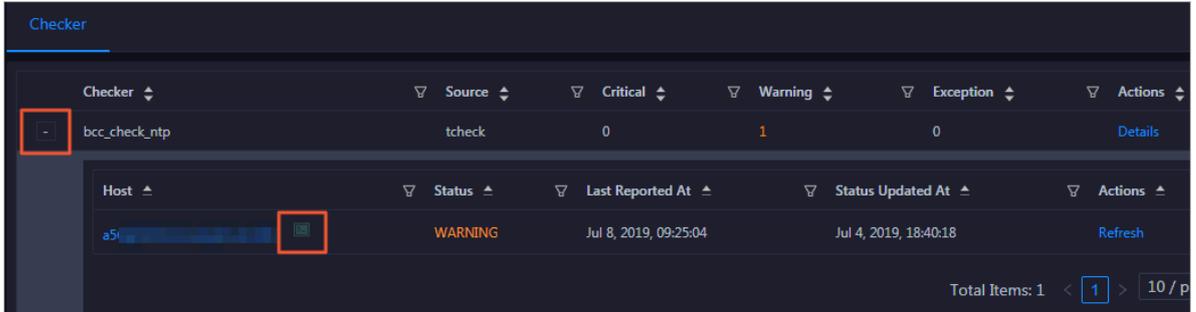
On the Health Status page, click Details in the Actions column of a checker with alerts. In the dialog box that appears, view the schemes to clear alerts.



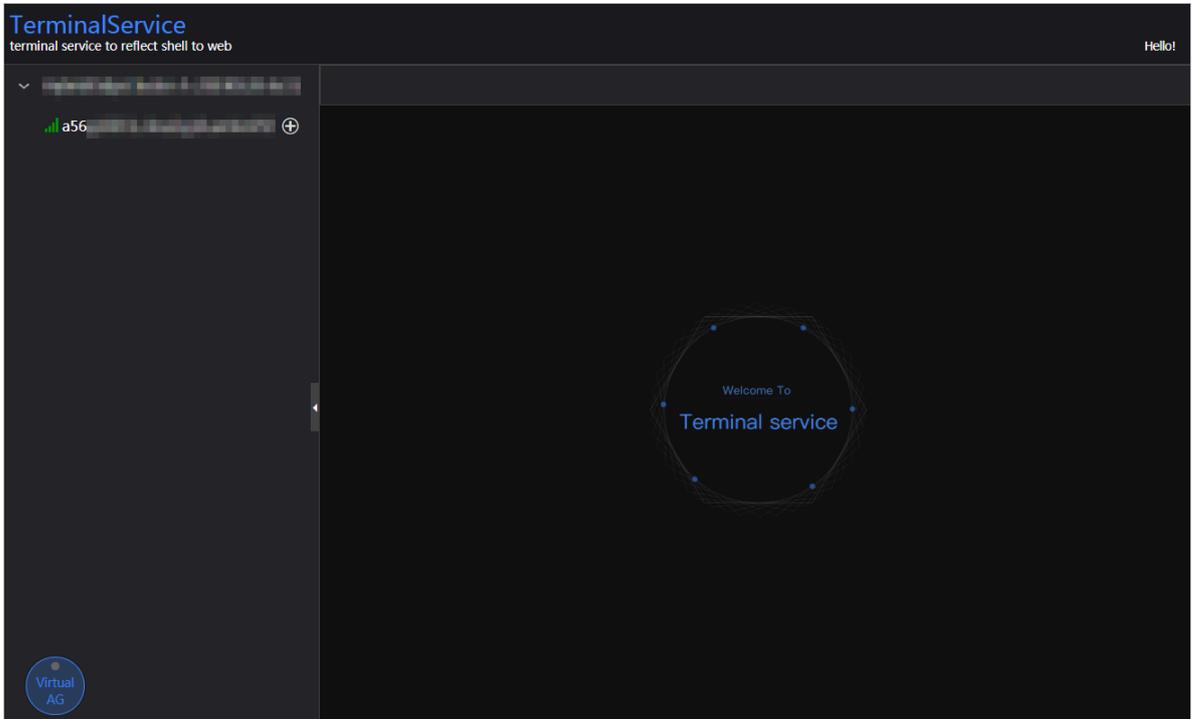
Log on to a host

To log on to a host to clear alerts or perform other operations, follow these steps:

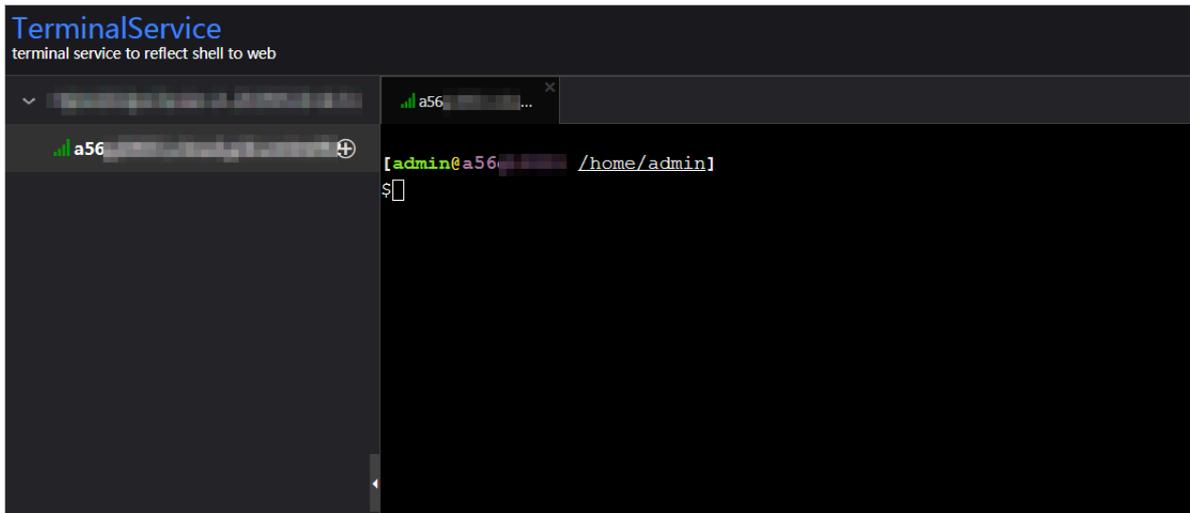
1. On the Health Status page, click + to expand a checker with alerts.



2. Click the Log On icon of a host. The TerminalService page appears.

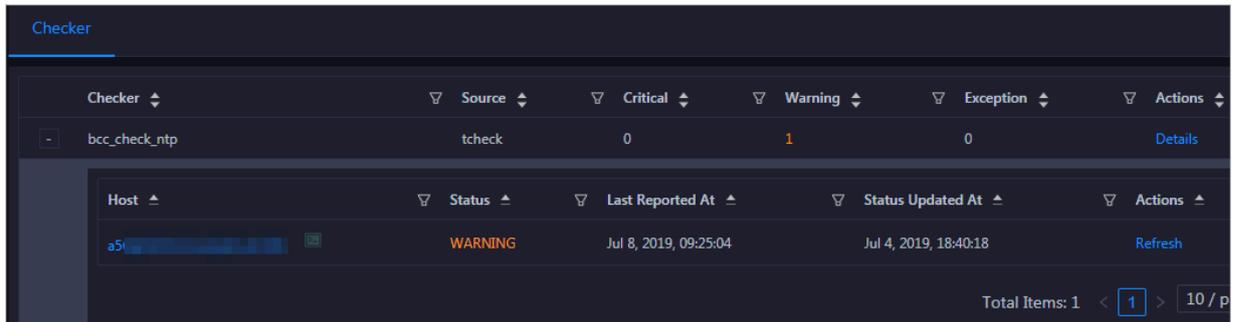


3. On the TerminalService page, click the hostname on the left to log on to the host.



Run a checker again

After you clear an alert for a host, click Refresh in the Actions column of the host to run the checker again for the host. In this way, you can check whether the alert is cleared.



3.9.2.5.4 Host services

On the Services page, you can view information about service instances and service instance roles of a host.

On the Hosts page, select a host in the left-side navigation pane, and then click the Services tab. The Services page for the host appears.

On the Services page, you can view the cluster, service instances, and service instance roles of the host.

3.9.3 Online O&M

3.9.3.1 Cluster health

You can view statistical information of Elasticsearch clusters. The cluster health information is the most important. An Elasticsearch cluster has three different health statuses: red, yellow, and green.

You can run the following command to check the health status of the cluster:

```
curl -u username:password http://domain:9200/_cluster/health
```

Color	Status	Description
Red	Some shards are unavailable.	The cluster contains unavailable shards , meaning that one or more indexes have unassigned shards.
Yellow	All shards are available , but some replicas are unavailable.	One or more indexes have unassigned replicas.
Green	All shards and replicas are available.	The cluster contains no unassigned shards or replicas.



Notice:

To ensure that your Elasticsearch cluster health status is green, all shards and replicas must be always available.

We recommend that `replica` (the number of replicas) is never greater than `amount_Node - 1` (`amount_Node` represents the number of nodes). This ensures that the health status of your Elasticsearch cluster is green after it is restarted when Dedicated Master is selected.

3.9.4 Troubleshooting

3.9.4.1 The cluster health status is yellow

If the health status of the cluster is yellow, operations such as changing the password or upgrading the cluster requires a long time. We recommend that you perform these operations when the health status of the cluster is green.

Cause: Some replicas are unassigned. You need to check which indexes in the cluster have unassigned replicas.

3.9.4.2 Query index status

You can run the following command to check which indexes have unassigned replicas.

```
curl -u username:password http://domain:9200/_cat/indices
```

If the cause of the error is that the number of replicas is greater than `amount_Node - 1`, you need to change the number of replicas of those indexes.

3.9.4.3 Restore index status

If your instance contains three nodes and an index has three replicas, the cluster health status is yellow. You can run the following command to set the number of replicas to 2:

```
curl -XPUT -u username:password http://domain:9200/index_name/_settings -H 'Content-Type: application/json' -d '{"index":{"number_of_replicas":(amount_Node - 1)}}'
```



Note:

After performing the operations (restart, scale-up, or customize configurations) on the instance, set a proper number of replicas based on the number of nodes. This improves the reliability and stability of the Elasticsearch cluster.

3.10 DataHub

3.10.1 Concepts and architecture

3.10.1.1 Terms

project

A project is an organizational unit in DataHub and contains one or more topics. DataHub projects and MaxCompute projects are independent of each other. Projects that you create in MaxCompute cannot be used in DataHub.

topic

The smallest unit for data subscription and publishing. You can use topics to distinguish different types of streaming data. For more information about projects and topics, see **Limits in *Product Introduction***.

time-to-live of records

The period that each record can be retained in the topic. Unit: day. Minimum value : 1. Maximum value: 7.

shard

A shard in a topic. Shards ensure the concurrent data transmission of a topic. Each shard has a unique ID. A shard can be in a different status. For more information about shard status, see the following table. Each active shard consumes server resources. We recommended that you create shards as needed.



Note:

Table 3-17: Shard status

Status	Description
Activating	All shards in a topic are in the Activating state when the topic is created. You cannot perform read or write operations on shards because they are being activated.
Active	Read and write operations are enabled when a shard is in the Active state.
Deactivating	A shard is in the Deactivating state when it is being split or merged with another shard. You cannot perform read or write operations on the shard because it is being deactivated.
Deactivated	A shard is in the Deactivated state when the split or merge operation is completed. The shard is read-only when it is in the Deactivated state.

hash key range

The range of hash key values for a shard, which is in [Starting hash key,Ending hash key) format. The hashing mechanism ensures that all records with the same partition key are written to the same shard.

merge

The operation that merges two adjacent shards. Two shards are considered adjacent if the hash key ranges for the two shards form a contiguous set with no gaps.

split

The operation that splits one shard into two adjacent shards.

record

A unit of data that is written into DataHub.

record type

The data type of records in a topic. Tuple and blob are supported. A tuple is a sequence of immutable objects. A blob is a chunk of binary data stored as a single entity.



Note:

- **The following data types are supported in a tuple topic.**

Table 3-18: Tuple data types

Type	Description	Value range
Bigint	An 8-byte signed integer. Note: Do not use the minimum value (-9223372036854775808) because this is a system reserved value.	-9223372036854775807 to 9223372036854775807
String	A string. Only UTF-8 encoding is supported.	The size of a string cannot exceed 1 MB.
Boolean	One of two possible values.	Valid values: True and False, true and false, or 0 and 1. <small>308 308</small>
Double	A double-precision floating-point number. It is 8 bytes in length.	-1.0 <small>10</small> to 1.0 <small>10</small>
Timestamp	A timestamp.	It is accurate to microseconds.

- **In a blob topic, a chunk of binary data is stored as a record. Records written into DataHub are Base64 encoded.**

Service roles

Table 3-19: Service roles

Product	Service role	Description
DataHub	Xstream	Receives read and write requests from the frontend server and forwards the requests to the Apsara Distributed File System.
	Shipper/Connector	Synchronizes data from DataHub to other Alibaba Cloud services, including MaxCompute, ApsaraDB RDS for MySQL, and Object Storage Service (OSS).
	Coordinator	Saves checkpoints for subscribed applications in the system. You can resume data consumption from any checkpoint you saved.
	Frontend	Receives all the read and write requests.

Run the following command on the admin gateway of a cluster to retrieve the services deployed on the cluster:

```
sr al
```

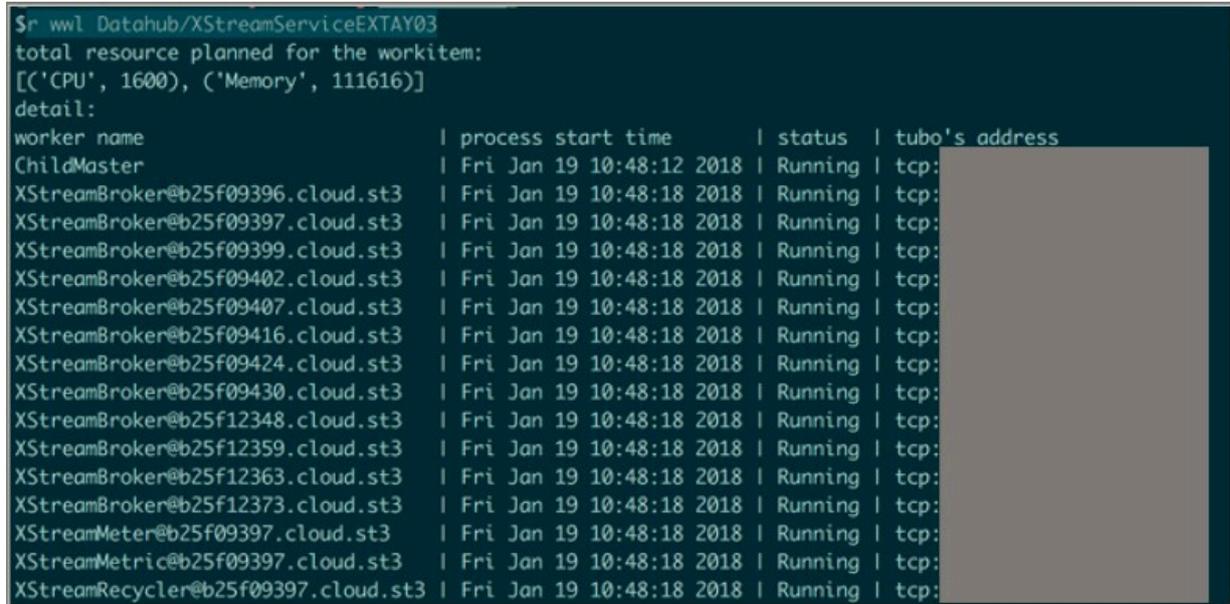
Figure 3-118: Services deployed on the cluster

```
[admin@datahub-ext-ay03-st3-ag /home/admin]
sr al
WorkItemName      | NuwaAddress
Datahub/ShipperServiceEXTAY03 | nuwa://datahub-ext-ay03-st3:10240/Datahub/ShipperServiceEXTAY03/ServiceMaster
Datahub/XStreamServiceEXTAY03 | nuwa://datahub-ext-ay03-st3:10240/Datahub/XStreamServiceEXTAY03/ServiceMaster
Datahub/CoordinatorServiceEXTAY03 | nuwa://datahub-ext-ay03-st3:10240/Datahub/CoordinatorServiceEXTAY03/ServiceMaster
```

Run the following command on the admin gateway of the cluster to retrieve the service role and servers where the service is running:

```
r wwl $WorkItemName
```

Figure 3-119: Service role and servers where the service is running

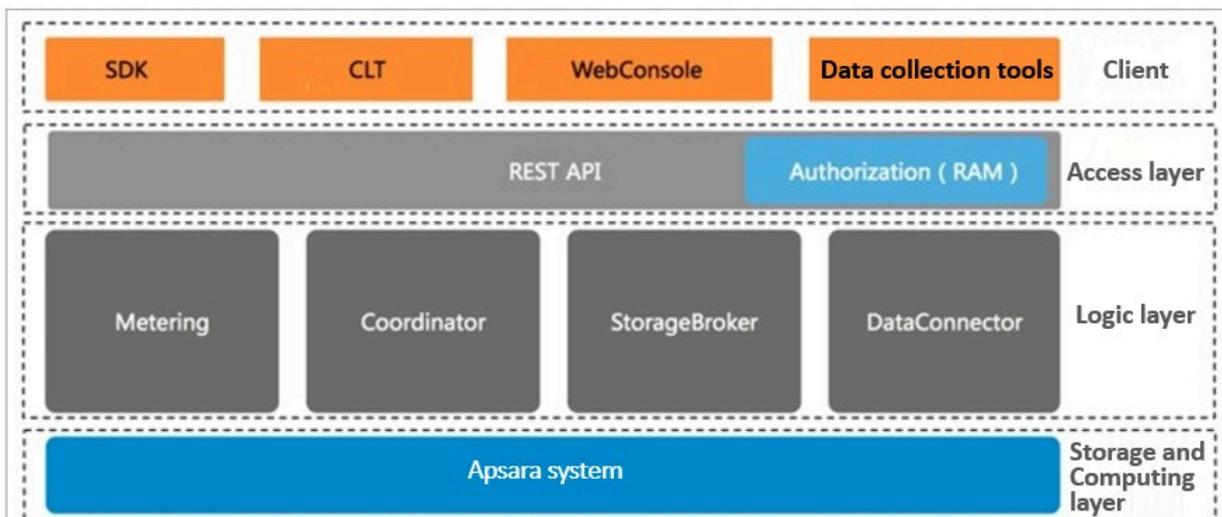


3.10.1.2 Architecture

3.10.1.2.1 Feature oriented architecture

Figure 3-120: Feature oriented architecture of DataHub shows the feature oriented architecture of DataHub.

Figure 3-120: Feature oriented architecture of DataHub



The architecture of DataHub consists of four layers: client, access layer, logical layer, and storage and scheduling layer.

Client

DataHub supports the following types of clients:

- **SDKs:** DataHub provides a variety of SDKs for C++, Java, Python, Ruby, and Go.
- **Command line tool (CLT):** You can run commands in Windows, Linux, or Mac operating systems to manage projects and topics.
- **Console:** In the console, you can manage projects and topics, create subscriptions, view shard details, monitor topic performance, and manage DataConnector.
- **Data collection tools:** Logstash, Fluentd, and Oracle GoldenGate (OGG).

Access layer

DataHub can be accessed through HTTP and HTTPS. DataHub supports RAM authorization and horizontal scaling of topic performance.

Logical layer

The logical layer handles the key features of DataHub, including project and topic management, data read and write, checkpoint-based data restoration, traffic statistics, and data archives. Based on these key features, the logical layer is composed of the following modules: StorageBroker, Metering, Coordinator, and DataConnector.

- **StorageBroker:** Enables the reading and writing of data in DataHub. Adopts the log file storage model of the Apsara Distributed File System, halving the read/write volume compared with the transfer of write-ahead logs. Stores three copies of data to ensure that no data is lost if a server fault occurs. Supports disaster recovery between data centers. Supports data write caching to ensure efficient consumption of real-time data. Supports independent read caching of historical data to enable concurrent consumption of the same data.
- **Metering:** Supports shard-level billing based on the consumption period.
- **Coordinator:** Supports checkpoint-based data restoration. Provides 150,000 QPS per node. Supports horizontal scaling of the processing capacity.
- **DataConnector:** Supports automatic data synchronization from DataHub to other Alibaba Cloud services, including MaxCompute, Object Storage Service (OSS), AnalyticDB, ApsaraDB RDS for MySQL, Table Store, and Elasticsearch.

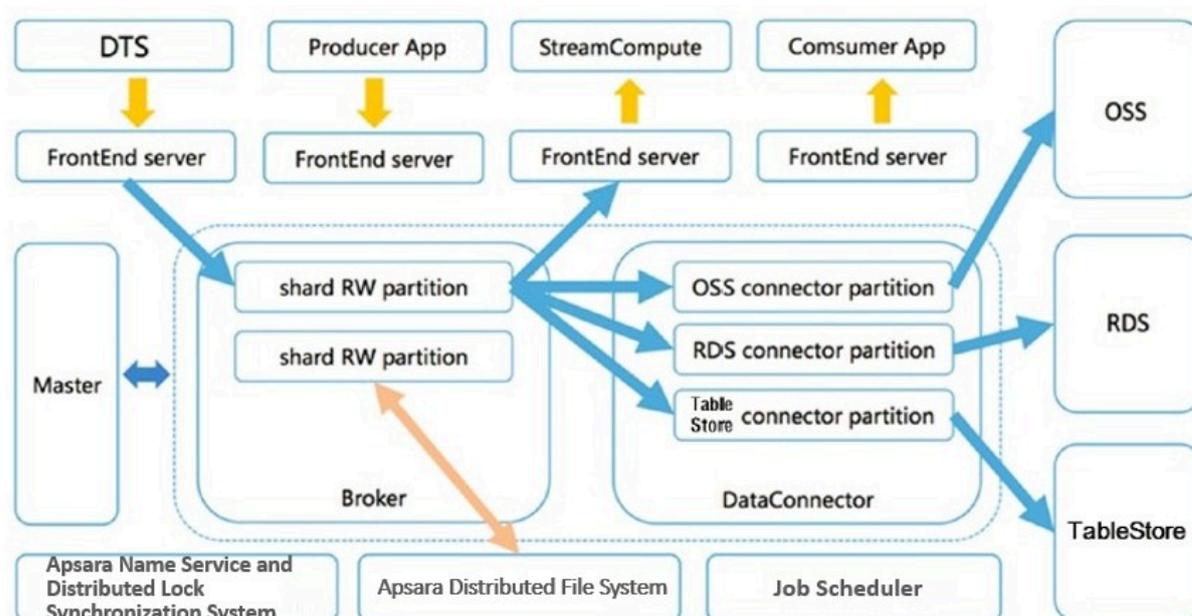
Storage and scheduling layer

- **Storage:** Based on the log file storage model of the Apsara Distributed File System , DataHub supports append operations and solid state drive (SSD) storage. Data in each shard is stored in a separate file based on the recording time of the data.
- **Scheduling:** Based on the scheduling module of Job Scheduler, DataHub assigns shards to nodes based on the traffic that occurs on each shard. This ensures that the shards do not occupy the CPU or memory of Job Scheduler. The number of partitions on a single node has no upper limit. DataHub supports failovers within milliseconds and hot upgrades.

3.10.1.2.2 Technical architecture

Figure 3-121: Technical architecture of DataHub shows the technical architecture of DataHub.

Figure 3-121: Technical architecture of DataHub



The figure shows the process from data ingestion to consumption.

1. A shard is the smallest unit of data management in DataHub, and is a first-in, first-out (FIFO) collection of records.
2. Data in each shard is stored in a set of log files on the Apsara Distributed File System.
3. The master distributes each shard to a broker. Each broker is responsible for the read and write operations of multiple shards.

4. The frontend server locates a broker based on the project, topic, and shard information specified in the request and forwards the request to the broker.
5. DataConnector reads data from the broker and forwards the data to other Alibaba Cloud services.

Data collector

You can write data into DataHub from applications developed by using SDKs and data collection tools such as LogStash, Fluentd, and Oracle GoldenGate. You can also write data by using Data Transmission Service (DTS) and Realtime Compute.

Frontend server

Frontend servers constitute the access layer and support horizontal scaling. You can call RESTful API operations to access DataHub. RAM authorization is supported

.

Master

The master handles metadata management and shard scheduling. It supports create, read, update, and delete operations on projects and topics. The master also supports split and merge operations on shards.

Broker

Brokers handle read and write operations on each shard including data indexing, caching, and file organization and management.

DataConnector

DataConnector forwards data in DataHub to other Alibaba Cloud services. DataConnector provides different features for various destination services. These features include automatically creating partitions in MaxCompute and converting data streams into files stored in OSS.

3.10.2 Commands and tools

3.10.2.1 Common commands for the Apsara system

DataHub is built based on the Apsara system. Both DataHub and the Apsara system including Job Scheduler, Apsara Distributed File System, and Apsara Name Service and Distributed Lock Synchronization System are hosted by Apsara Infrastructure Management Framework.

- **Run the following command to view the server roles that are installed on the server:**

```
tj_show
```

- **Run the following command to view all server roles:**

```
tj_show -l
```

- **Run the following command to retrieve a list of servers that the `pangu_chunkserver` server role is installed on:**

```
tj_show -r pangu.PanguChunkserver# //The hostnames of the servers  
are returned.  
tj_show -r pangu.PanguChunkserver# -ip //The IP addresses of the  
servers are returned.
```

- **Run the following command to retrieve a list of servers that the `FrontEnd` server role is installed on:**

```
tj_show -r datahub-frontend.Frontend#
```

- **Run the following command to retrieve a list of servers that the `WebConsole` server role is installed on:**

```
tj_show -r datahub-webconsole.WebConsole#
```

3.10.2.2 Common commands for Apsara Distributed File System

Commands for Apsara Distributed File System start with `pu` or `puadmin`. To view the complete description of a command, enter the command followed by `--help` and press enter.

- **Run the following command similar to the ls command used in Linux to retrieve the file content in a specific directory:**

```
pu ls
```

- **Run the following command to upload local files to Apsara Distributed File System:**

```
pu put
```

- **Run the following command to retrieve metadata:**

```
pu meta
```

- **Run the following command to retrieve details about all masters in Apsara Distributed File System:**

```
puadmin gems
```

- **Run the following command to retrieve details about all chunk servers:**

```
puadmin lscs
```

- **Run the following command to view version information:**

```
puadmin --buildinfo
```

- **Before maintaining a chunk server, remove the chunk server from the cluster. Perform the following operations:**

1. Run the following command to retrieve the current status of a chunk server:

```
pyadmin cs -stat tcp://x.x.x.x:10260
```

2. Run the following command to remove the chunk server from the cluster by setting its status to shutdown:

```
pyadmin cs -stat tcp://x.x.x.x:10260 --set=shutdown
```

3. After the maintenance is completed, run the following command to add the chunk server back to the cluster:

```
pyadmin cs -stat tcp://x.x.x.x:10260 --set=normal
```

3.10.2.3 Common commands for Job Scheduler

The commands for Job Scheduler start with `r`, which is encapsulation of `rpc.sh`.

```
alias r='sh /apsara/deploy/rpc_wrapper/rpc.sh'
```

- Run the following command to retrieve all services and service jobs:

```
r al
```



Note:

Typically, service jobs are deployed on the DataHub cluster. The list returned has many entries.

- Run the following command to retrieve the status of a service:

```
r wwl $servicename
```

- Run the following command to terminate a service:

```
r sstop $servicename
```

- Run the following command to start a service:

```
r sstart $servicename
```

- Run the following command to retrieve a list of all resources in the cluster:

```
r ttrl
```

- Run the following command to retrieve a list of idle resources in the cluster:

```
r tfrl
```

You can run other commands for scheduling purposes as needed.

3.10.2.4 Xstream

You can run commands on a service terminal by using Xstream for maintenance purposes.

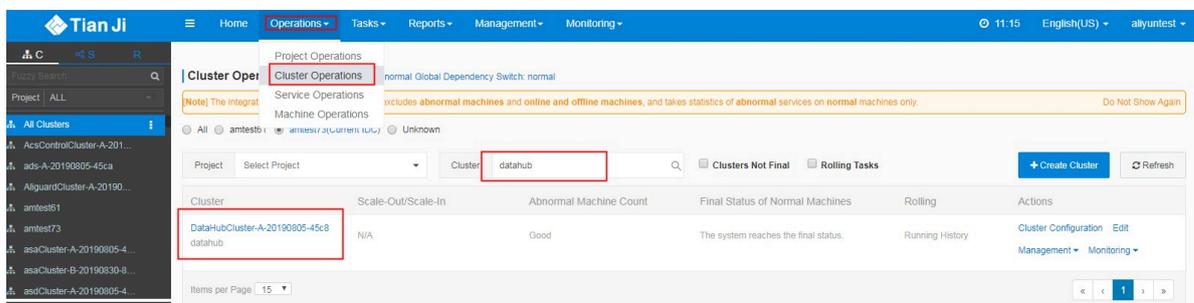


Note:

To use the Xstream tool, you must log on as the administrator.

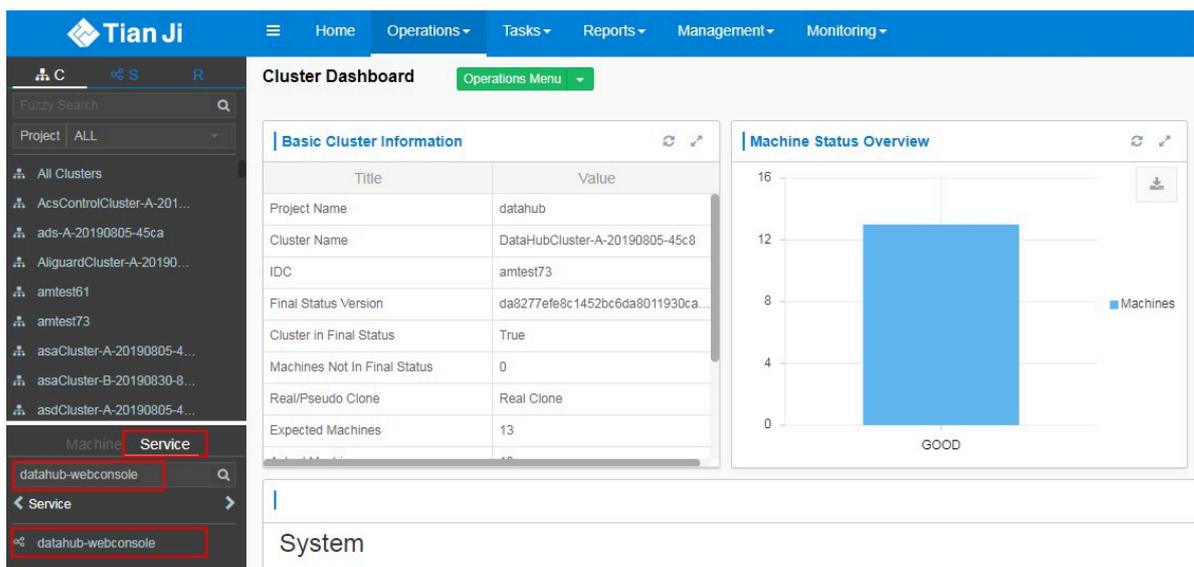
1. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, select **Operations > Cluster Operations**. In the left-side navigation pane, enter **datahub** in the search bar on the **C** tab page.

Figure 3-122: Search for a cluster



2. Click the cluster that appears in the search result. On the lower part of the left-side navigation pane, select the **Service** tab and search for **datahub-webconsole**.

Figure 3-123: Double-click the datahub-webconsole service



3. Double-click the datahub-webconsole service and double-click WebConsole#.

Figure 3-124: Double-click WebConsole#

**4. Hover your mouse pointer over the vertical dots next to the result that appears and select Terminal to open the TerminalService window.**

Figure 3-125: Click Terminal

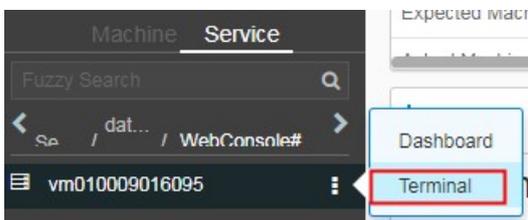
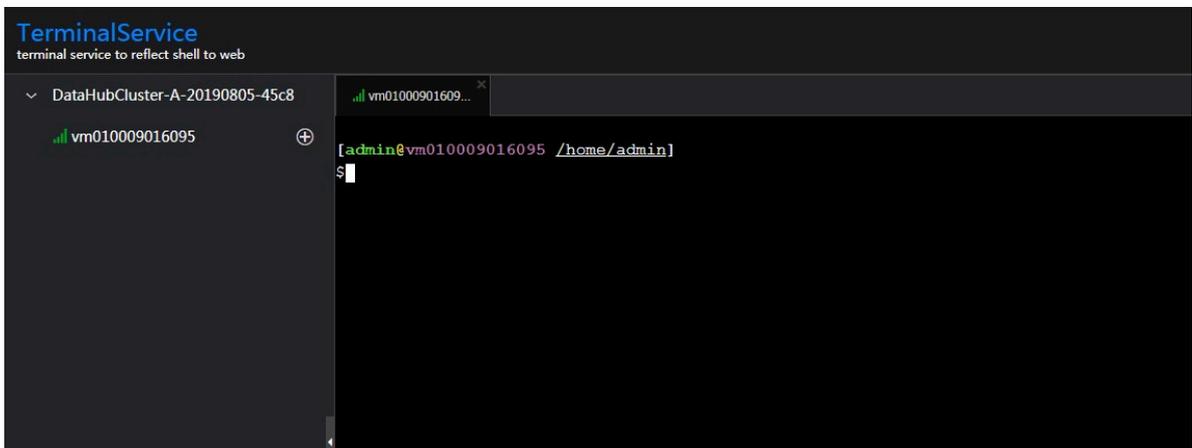


Figure 3-126: Open the TerminalService window

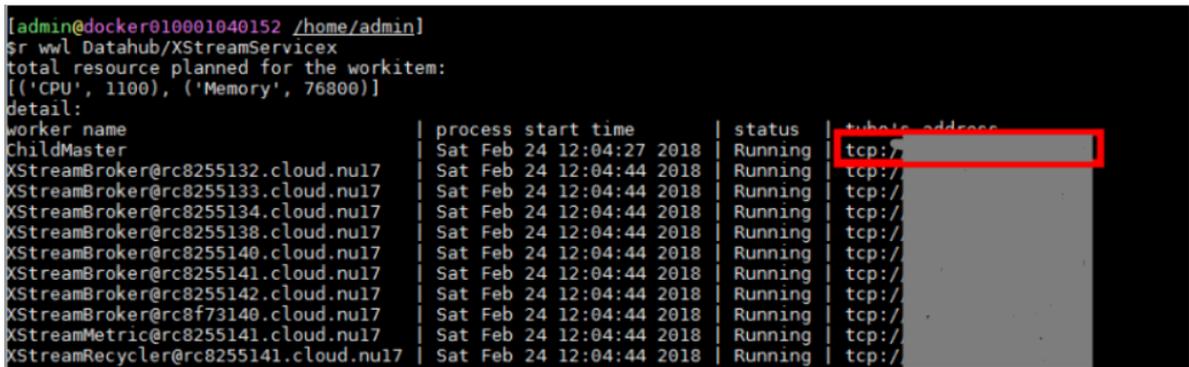


On the TerminalService window, you can use Xstream to run commands for operations and maintenance purposes.

1. Run the following command and find the IP address of the ChildMaster in the first line returned. Access the server that the ChildMaster is running by using Secure Shell (SSH).

```
r wwl Datahub/XStreamServiceX
```

Figure 3-127: Find the IP address of the ChildMaster



```
[admin@dockero10001040152 /home/admin]
$ r wwl Datahub/XStreamServiceX
total resource planned for the workitem:
[('CPU', 1100), ('Memory', 76800)]
detail:
worker name | process start time | status | tubob address
ChildMaster | Sat Feb 24 12:04:27 2018 | Running | tcp://...
XStreamBroker@rc8255132.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamBroker@rc8255133.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamBroker@rc8255134.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamBroker@rc8255138.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamBroker@rc8255140.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamBroker@rc8255141.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamBroker@rc8255142.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamBroker@rc8f73140.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamMetric@rc8255141.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
XStreamRecycler@rc8255141.cloud.nu17 | Sat Feb 24 12:04:44 2018 | Running | tcp://...
```

2. Run the following command to enter the specified directory:

```
cd /apsara/tubo/TempRoot/Datahub/XStreamServiceX/tool
```

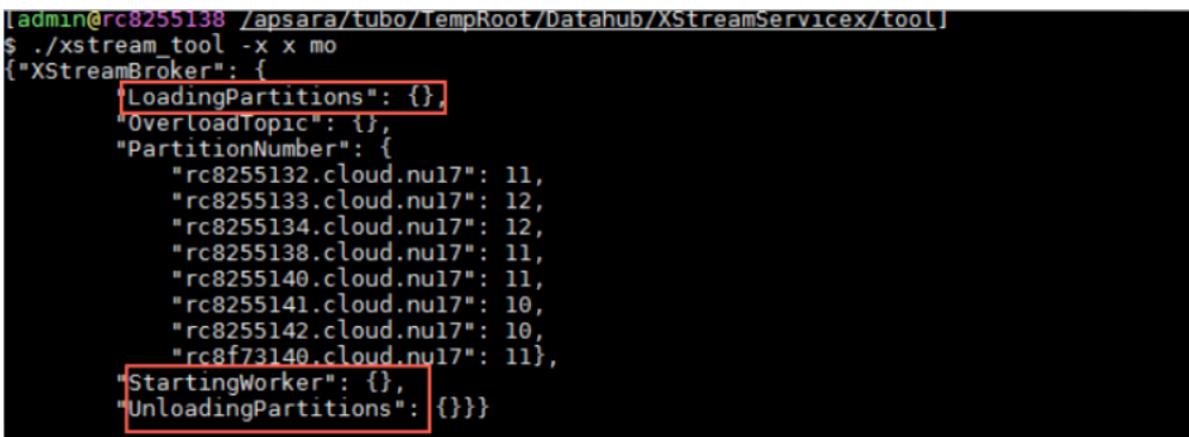
3. Run the following command to configure environment variables:

```
export LD_LIBRARY_PATH=/apsara/lib64/../../lib/
```

4. Run the following command to view resources:

```
./xstream_tool -x x mo
```

Figure 3-128: View resources



```
[admin@rc8255138 /apsara/tubo/TempRoot/Datahub/XStreamServiceX/tool]
$ ./xstream_tool -x x mo
{"XStreamBroker": {
  "LoadingPartitions": {},
  "OverloadTopic": {},
  "PartitionNumber": {
    "rc8255132.cloud.nu17": 11,
    "rc8255133.cloud.nu17": 12,
    "rc8255134.cloud.nu17": 12,
    "rc8255138.cloud.nu17": 11,
    "rc8255140.cloud.nu17": 11,
    "rc8255141.cloud.nu17": 10,
    "rc8255142.cloud.nu17": 10,
    "rc8f73140.cloud.nu17": 11},
  "StartingWorker": {},
  "UnloadingPartitions": {}}
```

If LoadingPartitions, UnloadingPartitions, and StartingWorker are returned with values, run the command again. If these parameters are repeatedly

returned with values, an error may occur when the shards are being activated or deactivated.

5. Run the following command to check the status of all brokers:

```
./xstream_tool gws -x x -r broker
```

Figure 3-129: Check the status of all brokers

```
admin@rc8255138 /apsara/tubo/tempRoot/Datahub/XStreamService/fool]
$ ./xstream_tool gws -x x -r broker
```

Machine Name	Requirement	Assignment	LoadedPartition	UnloadedPartition	UnconnectedWorker
rc8255132.cloud.nu17	1	1	11	0	0
rc8255133.cloud.nu17	1	1	12	0	0
rc8255134.cloud.nu17	1	1	12	0	0
rc8255138.cloud.nu17	1	1	11	0	0
rc8255140.cloud.nu17	1	1	11	0	0
rc8255141.cloud.nu17	1	1	10	0	0
rc8255142.cloud.nu17	1	1	10	0	0
rc8f73140.cloud.nu17	1	1	11	0	0
8	8	8	88	0	0

When 0 is returned for UnloadedPartition and UnconnectedWorker, the brokers are functioning properly.

6. Run the following command to check the status of all shards in the topic:

```
./xstream_tool -x x lsw -p $project -t $topic -r broker
```

Figure 3-130: Check the status of all shards in the topic

```
$ ./xstream_tool -x x lsw -p smoke_test_project -t datahub_to_datahub_input_1 -r broker
err_code: 0
err_msg: "Success"
workers {
  key: 3
  value: "Datahub/XStreamService/XStreamBroker@rc8255140.cloud.nu17"
}
workers {
  key: 5
  value: "Datahub/XStreamService/XStreamBroker@rc8255142.cloud.nu17"
}
workers {
  key: 2
  value: "Datahub/XStreamService/XStreamBroker@rc8255138.cloud.nu17"
}
workers {
  key: 7
  value: "Datahub/XStreamService/XStreamBroker@rc8255132.cloud.nu17"
}
workers {
  key: 4
  value: "Datahub/XStreamService/XStreamBroker@rc8255141.cloud.nu17"
}
```

From the command output, you can easily find the problematic shards.



Note:

We recommend that you do not run other commands by using Xstream except for those displayed in the examples. If you need to run other commands, contact an operations engineer.

3.10.3 Routine maintenance

3.10.3.1 Remove chunk servers from the DataHub cluster

Prerequisite

None.

Procedure

1. Check the environment.

a. Check the cluster status.

- Before performing the operation, check that the cluster is in final status and is running properly.

A. Log on to the Apsara Infrastructure Management Framework console. On the upper part of the left-side navigation pane, enter tianji in the search bar on the C tab page.

B. Hover over the vertical dots next to the cluster that appears in the search result. Then select Dashboard to go to the Cluster Dashboard page.

C. In the Service Instances List, check whether the Final Status of all instances for tianji is True.

D. Follow the same steps to navigate to the Cluster Dashboard page of the DataHub cluster and confirm that all DataHub instances are in final status. For assistance, contact an operations engineer.

- We recommend that you screenshot the key configuration information before performing the operation.

b. Note down the configurations of the servers to be removed from the cluster for verification purpose.



Note:

Do not remove any servers that are deployed with both frontend and chunk processes.

2. On the ops server, run the following commands to move one or more chunk servers from the DataHub cluster to the default cluster:

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
```

```
./tianji_ops_tool.py contract_nc -c clusterName -l machineList --  
config tianji_clt2.conf -s SRGname
```

**Note:****Parameters:**

- **-c: (Required)** The name of the cluster that the servers are removed from.
- **-s: (Required)** The name of the server role group (SRG) that the servers are removed from.
- **--config: (Required)** The tianji_clt configuration file.
- **-l: (Required)** The hostnames of the servers to be removed. Use commas (,) to separate multiple hostnames.

The command output is as follows:

```
machines removed from cluster XXX, newversion: 53111bf22f40e64d5343  
e0a73a2bce3241b7dac5  
-- Indicates that the operation is completed.
```

3. Check whether the servers are removed from the cluster.

- In the top navigation bar, select Operations > Cluster Operations. On the upper part of the left-side navigation pane, select the C tab and enter DataHub in the search bar.**
- Hover over the cluster that appears and select Cluster Operation and Maintenance Center to enter the Cluster Operation and Maintenance Center page.**
- Check whether the servers have been removed from the DataHub cluster.**

**Note:**

Click the number displayed next to Server s-out in the Scale-In Scale-Out box to check the corresponding server details. If the command is executed properly, the number displayed next to Server s-out gradually decreases until it turns 0.

4. Verify that the servers have been removed from the DataHub cluster.
 - a. Check whether the tianji and DataHub clusters are in final status.
 - A. Log on to the Apsara Infrastructure Management Framework console. On the upper part of the left-side navigation pane, select the C tab and enter tianji in the search bar.
 - B. Hover over the cluster that appears and click Dashboard to enter the Cluster Dashboard page.
 - C. In Service Instances List, check whether the Final Status of all instances for tianji is True.
 - D. Follow the same steps to navigate to the Cluster Dashboard page of the DataHub cluster and confirm that all DataHub instances are in final status. If you have any problems, contact an operations engineer.
 - b. Navigate to the Cluster Configuration page. Check whether the servers have been removed from the corresponding SRG by comparing the chunk server configurations against those noted down before the operation.

3.10.3.2 Add chunk servers to the DataHub cluster

Prerequisite

None.

Procedure

1. Check the environment.**a. Check the cluster status.**

- Before performing the operation, make sure that the cluster is in final status and is running properly.
 - A. Log on to the Apsara Infrastructure Management Framework console. On the upper part of the left-side navigation pane, enter tianji in the search bar on the C tab page.**
 - B. Hover over the vertical dots next to the cluster that appears in the search result. Then select Dashboard to go to the Cluster Dashboard page.**
 - C. In the Service Instances List, make sure that the Final Status of all instances for tianji is True.**
 - D. Follow the same steps to navigate to the Cluster Dashboard page of the DataHub cluster and confirm that all DataHub instances are in final status. If you have any problems, contact an operations engineer.**
- We recommend that you screenshot the key configuration information before performing the operation.
- Before adding chunk servers to the cluster, set the clone mode to normal. After the operation is completed, change the clone mode back to block. The path in the Apsara Infrastructure Management Framework console to

configure the clone mode is as follows: **Operations > Cluster Operations > Global Clone Switch.**

- Check whether the DataHub cluster is in Real Clone mode.

b. Check whether the server is added to Apsara Infrastructure Management Framework and used by the default cluster.

A. In the top navigation bar, select Operations > Machine Operations. On the Machine Operations page, check whether the server is in the default cluster.

B. If the server has not been added to Apsara Infrastructure Management Framework, upload a server configuration file to the framework. The procedure is as follows:

A. In the top navigation bar, select Operations > Machine Operations. On the Machine Operations page, click Machine Online/Offline.

B. In the dialog box that appears, click Download Schema on the Add Machine tab page to download the configuration file. Complete the information in the file. The nodename, Sn, and Ip parameters are required. Click Click to upload to upload the configuration file. After the file is uploaded, you can find the server in the default cluster.

2. On the ops server, run the following command to add a server to the specified cluster. After the command is executed, a rolling task begins.

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
./tianji_ops_tool.py expand_nc -c Cluster -s SRG -l machine1,
machine2 --config clt2.conf
```



Note:

Parameters:

- **-c: (Required)** The name of the cluster to be added with servers.
- **-s: (Required)** The name of the server role group (SRG) to be added with servers. The SRG must already exist in Apsara Stack Management Framework. You can check whether the SRG exists in the framework in the *machine_group.conf* file. If the SRG does not exist, deploy the SRG by using the Deployment Planner.
- **--config: (Required)** The tianji_clt configuration file.

- **-I:** (Required) The hostnames of the servers to be added. Use commas (,) to separate multiple hostnames.

3. Check the rolling progress in the Apsara Infrastructure Management Framework console.
4. Check the clone progress of the server in the Apsara Infrastructure Management Framework console.



Note:

In the top navigation bar, select Reports > System Reports. On the System Reports page, search for State of machine clone report.

5. When the rolling task is completed and the Server Status is GOOD, the server is added to the cluster.
6. Verify that the server has been added to the cluster.
 - a. Log on to the Apsara Infrastructure Management Framework console. Navigate to the Cluster Dashboard page of the DataHub cluster and check whether the cluster is in final status.
 - b. Check whether the server has been applied with the specified server role.
7. Check the monitoring data of the newly added server.

Log on to the Apsara Infrastructure Management Framework console. Navigate to the Server Details page of the newly added server and check whether there is monitoring data. If no monitoring data is available, restart the following three server roles: AcceleratorAgg#/AcceleratorMaster#/AcceleratorSource#.

8. Verify that the newly added server starts running properly.

Log on to the terminal of the newly added server. Run the `ps -ef | egrep "front | broker"` command to make sure that the server starts running properly.

Rollback plan

1. Remove the server that failed to be added to the DataHub cluster.

On the ops1 server, run the following commands to remove a server to the default cluster. If you need to remove multiple servers to the default cluster, run the commands to remove them one by one.

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
```

```
./tianji_ops_tool.py contract_nc -c Hybrid0dpsCluster-A-datahubx -s
ComputerServer -l machine1 --config clt2.conf
./tianji_ops_tool.py contract_nc -c Hybrid0dpsCluster-A-datahubx -s
ComputerServer -l machine2 --config clt2.conf
```

**Note:****Parameters:**

- **-c: (Required)** The name of the cluster that servers are removed from.
- **-s: (Required)** The name of the SRG that servers are removed from.
- **--config: (Required)** The tianji_clt configuration file.
- **-l: (Required)** The hostname of the server that failed to be added to the specified cluster.

2. Check whether the rollback operation takes effect.

- a. In the top navigation bar, select **Operations > Cluster Operations**. On the upper part of the left-side navigation pane, hover over the vertical dots next to the target cluster and click **Cluster Operation and Maintenance Center**.
- b. On the **Cluster Operation and Maintenance Center** page, make sure that the server has been removed.

3. Run the following commands to check the server removal task.

```
cd /cloud/app/tianji-tools/PrivateCloudTool#/tianji_ops_tool/current
ln -s /cloud/data/bootstrap_controller/BootstrapController#/
bootstrap_controller/tianji_dest.conf clt2.conf
./tianji_clt machinestatus -c default --config clt2.conf
```

4. Check whether the server information has been deleted from the cluster configuration file.

Navigate to the **Cluster Configuration File** page of the DataHub cluster and check whether the server information has been deleted from the cluster configuration file. If the server information is still in the `machine_group.conf` file, delete the information manually. Submit the changes to start a rolling task and wait for the task to complete.

3.10.3.3 Restore data after a power outage

Prerequisite

None.

Procedure

1. DataHub is based on the distributed storage technology of Apsara Distributed File System. Therefore, a power outage may cause data loss. After a power outage occurs, run the following command to check in the DataHub console whether data stored in Apsara Distributed File System has been lost.

```
puadmin fs -abnchunk|grep NONE|awk '{print $1}'|awk -F"_" '{print $1}'|while read line;do puadmin whois $line;done|grep FileId|awk '{print $4}' |sort|uniq >/home/admin/lostfile
-- Ignore directories that start with /deleted/ and send all other directories to an operations engineer to check for lost data.
```

2. Restore data based on file types.

- If DataHub files have been lost, inform your users to re-create the corresponding topics.
- If metadata has been lost, re-install the corresponding package or initialize the docker container.

3. After the data is restored, wait until the tianji cluster is in final status. For assistance, contact an operations engineer.

3.10.3.4 Reimage frontend and chunk servers

Prerequisite

If a system exception occurs and the hardware is functioning properly, reimage the server, server containers, and server roles.

Procedure

1. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, select **Operations > Machine Operations**. On the Machine Operations page, search for the physical server by hostname and click **Actions** to open the Operations Machine dialog box.
2. Select **Reclone** and click **Confirm**.
3. Check the clone task progress in the Apsara Infrastructure Management Framework console.

**Note:**

In the top navigation bar, select **Reports > System Reports**. On the System Reports page, you can find the State of machine clone report. In The Progress of Machine Clone, you can search for the server by machine name.

3.10.3.5 Upgrade or redeploy DataHub

Prerequisite

None.

Procedure

1. Check the environment.

- a. Log on to the Apsara Infrastructure Management Framework console. On the upper part of the left-side navigation pane, enter tianji in the search bar on the C tab page.**
- b. Hover over the vertical dots next to the cluster that appears in the search result and select Dashboard to go to the Cluster Dashboard page.**
- c. In the Service Instances List, check that the Final Status of all instances for tianji is True.**
- d. Follow the same steps to navigate to the Cluster Dashboard page of the DataHub cluster and confirm that all DataHub instances are in final status. If you have any problems, contact an operations engineer.**

2. Upgrade or redeploy DataHub.

- **Redeploy DataHub: If the DataHub service is unavailable due to exceptions on clusters, redeploy DataHub.**
 - a. Log on to the Apsara Infrastructure Management Framework console. In the top navigation bar, select Operations > Cluster Operations. On the Cluster Operations page, enter datahub in the search bar for Cluster.**
 - b. Click the cluster that appears in the search result to enter the Cluster Dashboard page. On the lower part of the navigation pane, click the Service**

tab, enter `datahub-webconsole` in the search bar, and then double-click the service.

- c. Double-click the `datahub-webconsole` service. Then double-click `WebConsole#`.
- d. On the `WebConsole#` tab page, hover over the vertical dots next to the result that appears and select `Terminal` to open the `TerminalService` window.
- e. Run the following command as the administrator to restart DataHub on all servers:

```
/home/admin/datahub_service/deploy/control start
```

- **Upgrade DataHub:** You can simultaneously upgrade the DataHub cluster and the output package from Apsara Infrastructure Management Framework console.
3. Verify that DataHub has been upgraded or re-deployed.
 - a. Log on to the Apsara Infrastructure Management Framework console. To search for the `tianji` cluster, in the left-side navigation pane, enter `tianji` in the search bar on the `C` tab page.
 - b. Hover over the vertical dots next to the `tianji` cluster and select `Dashboard` to go to the `Cluster Dashboard` page.
 - c. In the `Service Instances List`, check that the `Final Status` of all instances for `tianji` is `True`.
 - d. Follow the same steps to navigate to the `Cluster Dashboard` page of the DataHub cluster and confirm that all instances for DataHub are in final status. If you have any problems, contact an operations engineer.

3.10.3.6 Shut down problematic chunk servers

Prerequisite

None.

Procedure

1. Configure the action and action status for the problematic chunk server.

- a. Log on to the ops1 server and set action to rma and action status to pending for the problematic chunk server. In this example, the hostname of the problematic chunk server is m1.

Run the following command to configure the action and action status for the problematic chunk server:

```
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"
    }
  ]
}
```

**Note:**

Replace the IP address and hostname in the sample code with those of your problematic chunk server.

b. Configure audit logs.

```
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\n  \": \"/action/rma\", \n  \"description\" : \"/monitor/rma=error,\n  mtime: 1513488046851649\", \n  \"status\" : \"pending\""}'
```

**Note:**

- Replace the IP address and hostname in the sample code with those of your problematic chunk server.
- Replace the mtime parameter value in the sample code with the current time.

- **Run the following command to retrieve mtime. The sample code is for your reference only.**

The command is as follows:

```
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
  }
}
```

```
}

```

2. Wait for approval.

a. Check the action status of the server.

Run the following command to check the action status of the server:

```
curl "http://127.0.0.1:7070/api/v5/GetMachineInfo?hostname=m1"
```

The command output is a long list. We recommend that you search for the server by the key word "action_status": "pending".

After confirming that the action status is pending, you can approve the action in the Apsara Infrastructure Management Framework console.

A. In the top navigation bar, select **Operations > Cluster Operations**. On the upper part of the left-side navigation pane, enter **datahub** in the search bar on the **C** tab page.

B. Hover over the vertical dots next to the cluster that appears in the search result and select **Cluster Operation and Maintenance Center** to enter the **Cluster Operation and Maintenance Center** page.

C. In the **Server List of Cluster Operation and Maintenance Center**, find the chunk server to be shut down, click **Actions**, and then select **Approve Action**.

D. In the **Approve Action** dialog box, select **approved** for the action and click **OK**.

b. Check the action status of the server role. When the status is approved or done, you can shut down the server for maintenance.

Run the following command to check the action status:

```
curl http://127.0.0.1:7070/api/v5/GetMachineInfoPackage?hostname=m1&attr=sr.id,sr.action_name,sr.action_status
```

The command output is a long list of items. We recommend that you search for the server by the key word "action_status": "done".

3. After the action of the server changes to rma and action status changes to approved or done, shut down the server. Restart the server after the maintenance is completed.

4. After the server is restarted, run the following command to configure the action status of the server:

```
curl "http://127.0.0.1:7070/api/v5/SetMachineAction?hostname=m1&
action_name=rma" -d '{"action_name":"rma","action_status":"done", "
force":true}'
```

5. Log on to the Apsara Infrastructure Management Framework console. Navigate to the Cluster Dashboard page of the DataHub cluster to check whether the cluster is in final status.

3.10.3.7 Shut down the DataHub cluster

Prerequisite

None.

Procedure

1. Terminate DataHub services.

- a. On the webconsole server, run the following commands as an administrator and make sure that no data is returned:

```
puadmin fs -abnchunk -t none
puadmin fs -abnchunk -t onecopy
puadmin fs -abnchunk -t lessmin
```

- b. On the webconsole server, run the following commands as an administrator to terminate all services run by chunk servers in the Apsara system:

```
r ttrl |grep disk |awk '{print $1}' > tubo.list
pssh -h tubo.list -i "/apsara/cloud/tool/tianji/apsarad stop"
```

- c. On the webconsole server, run the following command as an administrator to make sure that all services in the Apsara system have been terminated:

```
pssh -h tubo.list -i "/apsara/cloud/tool/tianji/apsarad status"
```

2. Shut down the cluster.

3. Restart DataHub services.

- a. On the webconsole server, run the following command as an administrator to restart all services run by chunk servers in the Apsara system:

```
r ttrl |grep disk |awk '{print $1}' > tubo.list
```

```
pssh -h tubo.list -i "/apsara/cloud/tool/tianji/apsarad start"
```

- b. On the webconsole server, run the following command as an administrator to make sure that all services in the Apsara system are functioning properly:**

```
pssh -h tubo.list -i "/apsara/cloud/tool/tianji/apsarad status"
```

3.10.3.8 Replace a hard drive with a new one on the pangu_cs node

Prerequisite

Obtain the following information:

- **The hostname or the IP address.**
- **The drive letters of the problematic drive. For example, /dev/sdk.**
- **The ID of the problematic drive. For example, if the path of the problematic drive in the Apsara Distributed File System is /apsarapangu/disk5, the drive ID is 5. You can also obtain the drive ID by running the following command: puadmin lscs -m**

Procedure

- 1. Run the following command to check that the drive to be replaced is in DISK_ERROR status.**

```
puadmin lscs -m
```



Note:

If the hard drive is not in DISK_ERROR status, run the following command to change the status:

```
puadmin cs -stat tcp://hostname or IP address:10260 -d drive ID --set=ERROR
```

- 2. Run the following command to unmount the drive. In this example, the drive letters of the drive to be unmounted are /dev/sdk.**

```
sudo umount /dev/sdk1
```



Note:

Ignore this operation if the df command output shows that the drive is not mounted.

3. After the unmount operation is completed, replace the hard drive in hot swap mode.
4. Upload the `sudo repair_app_disk.sh` script to the server and execute the script to format the drive.
5. Run the following command to set the drive status in the Apsara Distributed File System to OK:

```
puadmin cs -stat tcp://hostname or IP address:10260 -d drive ID --set=OK
```

6. Restart the server. After the server is started up, it detects a new hard drive.



Note:

Kill the processes running on the `pangu_cs` chunk server and restart the server. Restarting a chunk server does not affect the continuity of your business because DataHub adopts a distributed storage model.

7. Run the following command to check whether the drive status is `DISK_OK`.

```
puadmin lscs -m
```

You can log on to the server to confirm that the drive has the `chunks` sub-directory. For example, the `chunks` exists in the `/apsarapangu/disk5/chunks/` directory and new chunks are written into the sub-directory.

3.10.4 Exceptions and solutions

This section describes some of the common error codes in the current version and corresponding solutions.

Error Code: `LimitEceeded`

Cause: The error code is returned because you can create up to 5 projects and 20 topics in a project in the previous version of DataHub.

Solution: In the latest version, you can create up to 10 projects and 1,000 topics in a project. Perform the following operations to change the project or topic limits:

1. Obtain the hostname of the ApsaraDB RDS for MySQL database from the following path: `/home/admin/datahub/service/deploy/env.cfg`.
2. Access the corresponding ApsaraDB RDS for MySQL database. In the `config_meta` table, check the values of `ProjectLimit4User` and `TopicLimit4Project`.

3. Run the following commands to update the configurations. The new configurations take about 1 minute to take effect. You do not need to restart the database.

```
update config_meta set config_value = 10 where config_type = 'ProjectLimit4User';
```

```
update config_meta set config_value = 1000 where config_type = 'TopicLimit4Project';
```

Error code: IanInvalidParameter

Cause: The error code is returned when StreamCompute attempts to capture records from DataHub by using an invalid timestamp. The timestamp you submit to the StreamCompute task is later than the current time, which may be caused by inaccurate local system time.

Solution: Correct your local system time by using the Network Time Protocol (NTP) or specify a timestamp that is for example 10 minutes earlier than the local system time.

Error code: InvalidCursor

Cause: The error code is returned when StreamCompute attempts to capture records from DataHub by using an invalid or expired cursor. An error may have occurred while StreamCompute is processing records from several days ago. When the time-to-live of the records expires and the records are deleted from DataHub, the cursor of these records is invalid.

Solution: Contact technical support for StreamCompute to learn about the cause of the task.

Error code: Parse response failed

Cause: This is probably caused by an invalid endpoint. For example, you may enter the console address as endpoint.

Solution: Perform a smoke test to check whether the system is running properly. If yes, check whether the endpoint is incorrect in the Apsara Infrastructure Management Framework console. Find the endpoint from the following path in the console: DataHubCluster > Cluster Dashboard > Cluster Resource > Service: datahub-frontend > dns in the Parameters and Result columns.

Error code: InternalServerError

Cause: Retry the smoke test or StreamCompute task. If the error code is still returned, an internal server error may occur. If the galaxy logs record this type of errors that occurred a long time ago, ignore these errors.

Solution: Use the following methods to search for corresponding logs to diagnose the issue. If you have any problems, screenshot the logs and contact technical support.

- In the logs directory of DataHubServer, search for the log files based on the specific time that the error occurred. The specific time can be found in the RequestId. RequestId is the unique ID of the request generated by DataHubServer.
- If more than one error occur, find the logs that are marked as ERROR in the logs directory of DataHubServer.

3.10.5 Appendix

3.10.5.1 Installation environment

Operation system: AliOS5U7-x86-64

Template: Bigdata

3.10.5.2 Deployment directories and services

Table 3-20: Services

Name	Type	Description
service-datahub-service	Controller	The service that is used to deploy DataHub backend services and used as the admin gateway of Apsara system.
service-datahub-webconsole	Controller	The service that is used to deploy the DataHub console and configured on the same container as service-datahub-service.
service-datahub-frontend	Worker	The service that is used to deploy frontend servers and used as chunk servers.

Name	Type	Description
Chunkserver	Worker	The service that is used to deploy chunk servers in Apsara Distributed File System.
PanguMaster	Controller	The service that is used to deploy three masters in Apsara Distributed File System.
NuwaMaster	Controller	The service that is used to deploy three masters of Apsara Name Service and Distributed Lock Synchronization System.
FuxiMaster	Controller	The service that is used to deploy two masters of Job Scheduler.

Table 3-21: Deployment directories and corresponding services

Module	Directory	Service
Datahub/XStreamService	/home/admin/datahub_service	service-datahub-service
Datahub/ShipperService	/home/admin/datahub_service	service-datahub-service
Datahub/CoordinatorService	/home/admin/datahub_service	service-datahub-service
WebConsole	/home/admin/datahub_webconsole	service-datahub-webconsole
Smoke	/home/admin/datahub_smoke	service-datahub-frontend
Frontend	/home/admin/datahub_frontend_server	service-datahub-frontend

3.10.5.3 Error codes

Table 3-22: Error codes

Error code	HTTP status code	Description
InvalidUriSpec	400	The error code is returned when the request URI is invalid. This is probably caused by invalid topic or project names.
InvalidParameter	400	The error code is returned when a parameter is invalid. For more information about the cause of the error, see the error message.
Unauthorized	401	The error code is returned when the signature is incorrect. This is usually caused by an incorrect AccessKey or a time difference of more than 15 minutes between the client and the server.
NoPermission	403	The error code is returned when you do not have the permission to perform the operation.
InvalidSchema	400	The error code is returned when the schema format is invalid.
InvalidCursor	400	The error code is returned when the cursor is invalid or has expired.
NoSuchProject	404	The error code is returned when the specified project does not exist.
NoSuchTopic	404	The error code is returned when the specified topic does not exist.
NoSuchShard	404	The error code is returned when the specified shard ID does not exist.
ProjectAlreadyExist	400	The error code is returned when the project name already exists.
TopicAlreadyExist	400	The error code is returned when the topic name already exists.

Error code	HTTP status code	Description
InvalidShardOperation	405	The error code is returned when the operation on the shard is not allowed . For example, you are not allowed to write data into a shard when it is in Deactivated status.
LimitExceeded	400	The error code is returned when a specified threshold is exceeded. For example, you create no more than 512 shards in a topic and 20 topics in a project.
InternalServerError	500	The error code is returned when an unknown or internal error occurs or when the system is being upgraded. For more information about the cause of the error, obtain the request ID or search DataHub server logs for InternalServerError.