

MapReduce Service

User Guide

Issue 01
Date 2024-09-23



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 Preparations.....	1
1.1 Configuring MRS Cloud Service Authorization.....	1
1.2 Creating an IAM User and Granting MRS Permissions.....	2
1.3 Creating a Custom Policy for MRS.....	7
2 MRS Cluster Planning.....	13
2.1 Service Selection.....	13
2.1.1 MRS Cluster Types.....	13
2.1.2 MRS Cluster Node Types.....	14
2.1.3 MRS Cluster Node Specifications.....	16
2.2 MRS Cluster Deployment.....	17
2.2.1 Overview.....	17
2.2.2 Kerberos Authentication for MRS Clusters.....	25
2.2.3 ECS Specifications Supported by MRS Clusters.....	27
3 Buying MRS Clusters.....	30
3.1 Quickly Buying an MRS Cluster.....	30
3.2 Manually Buying an MRS Cluster.....	34
4 Installing an MRS Cluster Client.....	49
4.1 Installing a Client (MRS 3.x).....	49
4.2 Installing a Client (MRS 2.x or Earlier).....	59
5 Submitting an MRS Job.....	65
5.1 MRS Job Types.....	65
5.2 Uploading Application Data to an MRS Cluster.....	66
5.3 Running an MRS Job.....	69
5.3.1 Running a MapReduce Job.....	69
5.3.2 Running a SparkSubmit Job.....	73
5.3.3 Running a HiveSQL Job.....	80
5.3.4 Running a Spark SQL Job.....	84
5.3.5 Running a Flink Job.....	89
5.3.6 Running a HadoopStreaming Job.....	96
5.4 Viewing MRS Job Details and Logs.....	98
6 Managing Clusters.....	101

6.1 Overview.....	101
6.2 Introduction to MRS Manager.....	103
6.3 Accessing MRS Manager.....	108
6.4 Managing an MRS Cluster.....	115
6.4.1 Viewing Basic Information About an MRS Cluster.....	115
6.4.2 Checking the Running Status of an MRS Cluster.....	118
6.4.3 Starting and Stopping an MRS Cluster.....	121
6.4.4 Restarting an MRS Cluster.....	122
6.4.5 Exporting MRS Cluster Configuration Parameters.....	125
6.4.6 Synchronizing the MRS Cluster Configuration.....	126
6.4.7 Transforming a Pay-per-Use MRS Cluster to a Yearly/Monthly Cluster.....	128
6.4.8 Deleting an MRS Cluster.....	128
6.4.9 Changing the VPC Subnet of an MRS Cluster.....	129
6.4.10 Replacing the NTP Server for an MRS Cluster.....	132
6.4.11 Modifying the OMS Service Configuration.....	134
6.4.12 Modifying MRS Manager Routing Table.....	136
6.5 Managing MRS Cluster Components.....	139
6.5.1 Checking the Running Status of an MRS Cluster Component.....	139
6.5.2 Starting and Stopping an MRS Cluster Component.....	146
6.5.3 Restarting an MRS Cluster Component.....	147
6.5.4 Adding and Deleting an MRS Cluster Component.....	154
6.5.5 Modifying the Configuration Parameters of an MRS Cluster Component.....	155
6.5.6 Viewing the Modified Component Configuration Parameters of an MRS Cluster.....	160
6.5.7 Synchronizing MRS Component Configuration Parameters.....	162
6.5.8 Adding Custom MRS Component Parameters.....	166
6.5.9 Managing MRS Role Instances.....	170
6.5.10 Managing MRS Role Instance Groups.....	172
6.5.11 Modifying MRS Role Instance Parameters.....	174
6.5.12 Perform an Active/Standby Switchover for MRS Role Instances.....	176
6.5.13 Decommissioning and Recommissioning an MRS Role Instance.....	177
6.5.14 Enabling and Disabling Ranger Authentication for an MRS Component.....	179
6.5.15 Accessing Web Pages of Open Source Components Managed in MRS Clusters.....	181
6.6 Managing MRS Cluster Nodes.....	182
6.6.1 Checking the Running Status of an MRS Cluster Node.....	182
6.6.2 Starting and Stopping All Roles on an MRS Cluster Node.....	186
6.6.3 Isolating an MRS Cluster Node.....	187
6.6.4 Modifying the Rack Information of an MRS Cluster Node.....	189
6.6.5 Scaling Up Master Node Specifications in an MRS Cluster.....	192
6.6.6 Synchronizing Disk Information of an MRS Cluster Node.....	198
6.6.7 Adding a Tag to an MRS Cluster/Node.....	199
6.6.8 Configuring Bootstrap Actions for an MRS Cluster Node.....	203
6.6.8.1 MRS Bootstrap Action Overview.....	203

6.6.8.2 Preparing the Bootstrap Action Script for an MRS Node.....	203
6.6.8.3 Adding MRS Node Bootstrap Actions and Installing Third-Party Software.....	204
6.6.8.4 Viewing the Bootstrap Action Execution Records of an MRS Node.....	208
6.7 Managing the MRS Cluster Client.....	209
6.7.1 Updating the MRS Cluster Client After the Server Configuration Expires.....	209
6.7.2 Viewing the Installed MRS Cluster Client.....	213
6.7.3 Batch Upgrading MRS Cluster Clients.....	214
6.8 Managing MRS Cluster Jobs.....	216
6.8.1 Stopping and Deleting an MRS Cluster Job.....	216
6.8.2 Configuring Notification Rules for MRS Jobs.....	217
6.9 Managing MRS Cluster Tenants.....	218
6.9.1 Introduction to MRS Multi-Tenancy.....	218
6.9.2 Using MRS Multi-Tenancy.....	230
6.9.3 Configuring MRS Tenants.....	232
6.9.3.1 Creating an MRS Tenant.....	233
6.9.3.2 Creating an MRS Sub-Tenant.....	246
6.9.3.3 Binding Tenant to an MRS Cluster User.....	261
6.9.3.4 Adding an MRS Tenant Resource Pool.....	264
6.9.3.5 Configuring the Queue Capacity Policy of a Resource Pool.....	266
6.9.3.6 Configuring the MRS Tenant Queue.....	269
6.9.4 Managing MRS Tenant Resources.....	277
6.9.4.1 Managing the MRS Tenant Resource Directory.....	277
6.9.4.2 Managing MRS Tenant Resource Pools.....	279
6.9.4.3 Clearing the MRS Tenant Queue Configuration.....	281
6.9.4.4 Restoring MRS Tenant Data After YARN Is Reinstalled.....	282
6.9.4.5 Deleting an MRS Tenant.....	282
6.9.4.6 Managing Global User Policies When Using Superior Scheduler.....	284
6.9.4.7 Clearing Tenant's Non-Associated Queues Using Capacity Scheduler.....	286
6.9.5 Switching the MRS Tenant Resource Scheduler.....	287
6.10 Managing MRS Cluster Users.....	290
6.10.1 Cluster User Permissions.....	290
6.10.1.1 MRS Cluster User Permission Model.....	291
6.10.1.2 MRS Cluster User Identity Authentication Policy.....	293
6.10.1.3 MRS Cluster User Permission Authentication Policy.....	295
6.10.1.4 Default Permissions of the MRS Cluster.....	297
6.10.1.5 Synchronizing IAM Users to MRS.....	300
6.10.2 MRS Cluster User Accounts.....	305
6.10.3 Managing MRS Cluster Roles.....	359
6.10.4 Managing MRS Cluster User Groups.....	362
6.10.5 Managing MRS Cluster Users.....	364
6.10.5.1 Creating an MRS Cluster User.....	365
6.10.5.2 Modifying MRS Cluster User Information.....	367

6.10.5.3 Locking an MRS Cluster User.....	368
6.10.5.4 Deleting an MRS Cluster User.....	370
6.10.5.5 Initializing MRS Cluster User Passwords.....	372
6.10.5.6 Downloading MRS Cluster User Credentials.....	373
6.10.6 Unlocking an MRS Cluster User.....	374
6.10.6.1 Unlocking an LDAP User in the MRS Cluster.....	374
6.10.6.2 Unlocking the LDAP Management Account of the MRS Cluster.....	377
6.10.7 Configuring Password Policies for MRS Cluster Users.....	377
6.10.8 Configuring the Private Attribute of MRS Cluster Users.....	381
6.11 Managing MRS Cluster Metadata.....	382
6.11.1 MRS Cluster Metadata Overview.....	382
6.11.2 Storing Ranger Metadata to RDS.....	383
6.11.3 Storing Hive Metadata to RDS.....	393
6.11.4 Configuring a LakeFormation Data Connection.....	398
6.11.4.1 LakeFormation Overview.....	399
6.11.4.2 Preparing for a LakeFormation Data Connection.....	400
6.11.4.3 Configuring a LakeFormation Data Connection During Cluster Creation.....	409
6.11.5 Managing MRS Cluster Data Connections.....	415
6.12 Managing Static Service Resources in an MRS Cluster.....	416
6.12.1 Overview of Static Service Resources.....	416
6.12.2 Configuring Static Resources for an MRS Cluster.....	417
6.12.3 Checking the Static Resources of an MRS Cluster.....	423
7 MRS Cluster O&M.....	426
7.1 Cluster O&M.....	426
7.2 Logging In to an MRS Cluster.....	430
7.2.1 Checking MRS Active/Standby Management Nodes.....	430
7.2.2 Logging In to an MRS Cluster Node.....	431
7.3 Viewing MRS Cluster Monitoring Metrics.....	436
7.3.1 Viewing MRS Cluster Resource Monitoring Metrics.....	437
7.3.2 Viewing MRS Cluster Component Monitoring Metrics.....	440
7.3.3 Viewing MRS Node Resource Monitoring Metrics.....	446
7.3.4 Dumping MRS Cluster Monitoring Data.....	451
7.4 Checking MRS Cluster Health.....	454
7.4.1 Performing a Health Check for an MRS Cluster.....	455
7.4.2 Performing Health Checks on MRS Cluster Nodes.....	456
7.4.3 Viewing and Exporting a Health Check Report.....	457
7.5 Adjusting the Capacity of an MRS Cluster.....	459
7.5.1 Scaling Out an MRS Cluster.....	459
7.5.2 Expanding a Data Disk of an MRS Cluster Node.....	464
7.5.3 Scaling In an MRS Cluster.....	466
7.5.4 Scaling In ClickHouseServer Nodes.....	471
7.5.5 Unsubscribing from a Specified Node in a Yearly/Monthly MRS Cluster.....	478

7.5.6 MRS Task Node Auto Scaling.....	480
7.5.6.1 Automatic Scaling of Task Nodes in an MRS Cluster.....	480
7.5.6.2 Adding an Auto Scaling Policy for MRS Task Nodes.....	488
7.5.6.3 Managing MRS Cluster Auto Scaling Policies.....	494
7.6 MRS Cluster Data Backup and Restoration.....	496
7.6.1 Backing Up and Restoring MRS Cluster Data.....	497
7.6.2 Enabling MRS Inter-Cluster Replication.....	504
7.6.3 Creating an MRS Cluster Data Backup Task.....	506
7.6.4 Creating an MRS Cluster Data Restoration Task.....	508
7.6.5 Backing Up MRS Cluster Component Data.....	509
7.6.5.1 Backing Up Manager Data (MRS 2.x and Earlier).....	509
7.6.5.2 Backing Up Manager Data (MRS 3.x and Later Versions).....	510
7.6.5.3 Backing Up CDL Service Data.....	515
7.6.5.4 Backing Up ClickHouse Metadata.....	517
7.6.5.5 Backing Up ClickHouse Service Data.....	520
7.6.5.6 Backing Up DBService Data.....	525
7.6.5.7 Backing Up Flink Metadata.....	529
7.6.5.8 Backing Up HBase Metadata.....	532
7.6.5.9 Backing Up HBase Service Data.....	535
7.6.5.10 Backing Up HDFS NameNode Data.....	540
7.6.5.11 Backing Up HDFS Service Data.....	544
7.6.5.12 Backing Up Hive Service Data.....	549
7.6.5.13 Backing Up IoTDB Metadata.....	555
7.6.5.14 Backing Up IoTDB Service Data.....	558
7.6.5.15 Backing Up Kafka Metadata.....	560
7.6.6 Restoring MRS Cluster Component Data.....	564
7.6.6.1 Restoring Manager Data (MRS2.x and Earlier).....	564
7.6.6.2 Restoring Manager Data (MRS 3.x and Later Versions).....	567
7.6.6.3 Restoring CDL Service Data.....	572
7.6.6.4 Restoring ClickHouse Metadata.....	574
7.6.6.5 Restoring ClickHouse Service Data.....	576
7.6.6.6 Restoring DBService Metadata.....	579
7.6.6.7 Restoring Flink Metadata.....	583
7.6.6.8 Restoring HBase Metadata.....	585
7.6.6.9 Restoring HBase Service Data.....	589
7.6.6.10 Restoring HDFS NameNode Metadata.....	593
7.6.6.11 Restoring HDFS Service Data.....	597
7.6.6.12 Restoring Hive Service Data.....	601
7.6.6.13 Restoring IoTDB Metadata.....	606
7.6.6.14 Restoring IoTDB Service Data.....	609
7.6.6.15 Restoring Kafka Metadata.....	611
7.6.7 Managing MRS Cluster Backup and Restoration Tasks.....	615

7.6.8 Using HDFS Snapshots to Quickly Restore Component Service Data.....	618
7.7 MRS Cluster Patching.....	619
7.7.1 Viewing Patch Information for an MRS Cluster.....	619
7.7.2 Patching an MRS Cluster.....	620
7.7.3 Applying Rolling Patches for an MRS Cluster.....	621
7.7.4 Patching Hosts Isolated in an MRS Cluster.....	625
7.8 MRS Cluster Patch Description.....	626
7.8.1 MRS 3.0.5.1 Patch Description.....	626
7.8.2 MRS 2.1.0.11 Patch Description.....	628
7.8.3 MRS 2.1.0.10 Patch Description.....	634
7.8.4 MRS 2.1.0.9 Patch Description.....	639
7.8.5 MRS 2.1.0.8 Patch Description.....	643
7.8.6 MRS 2.1.0.7 Patch Description.....	647
7.8.7 MRS 2.1.0.6 Patch Description.....	650
7.8.8 MRS 2.1.0.3 Patch Description.....	653
7.8.9 MRS 2.1.0.2 Patch Description.....	654
7.8.10 MRS 2.1.0.1 Patch Description.....	656
7.8.11 MRS 2.0.6.1 Patch Description.....	657
7.8.12 MRS 2.0.1.3 Patch Description.....	658
7.8.13 MRS 2.0.1.2 Patch Description.....	659
7.8.14 MRS 2.0.1.1 Patch Description.....	660
7.8.15 MRS 1.9.3.3 Patch Description.....	660
7.8.16 MRS 1.9.3.1 Patch Description.....	662
7.8.17 MRS 1.9.2.2 Patch Description.....	663
7.8.18 MRS 1.9.0.8, 1.9.0.9, and 1.9.0.10 Patch Description.....	664
7.8.19 MRS 1.9.0.7 Patch Description.....	669
7.8.20 MRS 1.9.0.6 Patch Description.....	673
7.8.21 MRS 1.9.0.5 Patch Description.....	676
7.8.22 MRS 1.8.10.1 Patch Description.....	679
7.9 Viewing Logs of an MRS Cluster.....	679
7.9.1 Overview of MRS Cluster Logs.....	679
7.9.2 Viewing MRS Operation Logs.....	708
7.9.3 Viewing MRS Cluster History.....	709
7.9.4 Viewing MRS Cluster Audit Logs.....	711
7.9.5 Viewing Role Instance Logs of MRS Components.....	713
7.9.6 Searching for MRS Cluster Logs Online.....	714
7.9.7 Downloading MRS Cluster Logs.....	717
7.9.8 Collecting MRS Cluster Service Stack Information.....	718
7.9.9 Configuring Default Log Level and Archive File Size for MRS Components.....	721
7.9.10 Configuring the Number of Local Backups of MRS Cluster Audit Logs.....	722
7.9.11 Configuring Dumping for MRS Cluster Audit Logs.....	723
7.10 MRS Cluster Security Configuration.....	727

7.10.1 Cluster Mutual Trust Management.....	727
7.10.1.1 Overview of Mutual Trust Between MRS Clusters.....	727
7.10.1.2 Changing the System Domain Name of an MRS Cluster.....	728
7.10.1.3 Configuring Mutual Trust Between MRS Clusters.....	732
7.10.1.4 Configuring User Permissions for Mutually Trusted MRS Clusters.....	739
7.10.2 Replacing MRS Cluster Certificates.....	741
7.10.2.1 Replacing the CA Certificate.....	741
7.10.2.2 Replacing an HA Certificate.....	744
7.10.3 MRS Cluster Security Hardening.....	747
7.10.3.1 MRS Cluster Security Hardening Policies.....	747
7.10.3.2 Configuring Hadoop Data Encryption During Transmission.....	749
7.10.3.3 Configuring Kafka Data Encryption During Transmission.....	752
7.10.3.4 Configuring HDFS Data Encryption During Transmission.....	753
7.10.3.5 Configuring Spark Data Encryption During Transmission.....	756
7.10.3.6 Configuring ZooKeeper Data Encryption During Transmission.....	757
7.10.3.7 Encrypting Data Transmission Between the Controller and Agent.....	758
7.10.3.8 Configuring a Trusted IP Address to Access LDAP.....	760
7.10.3.9 HFile and WAL Encryption.....	762
7.10.3.10 Configuring the IP Address Whitelist for Modifying Data in an HBase Read-Only Cluster.....	769
7.10.3.11 Configuring LDAP Output Audit Logs.....	770
7.10.3.12 Updating Encryption Keys of an MRS Cluster.....	771
7.10.3.13 Updating the SSH Key of User omm on MRS Cluster Nodes.....	773
7.10.3.14 Enabling and Disabling Permission Verification on MRS Cluster Components.....	775
7.10.3.15 Allowing External Users to Access MRS Clusters in Normal Mode.....	777
7.10.3.16 Configuring Secure Communication Authorization for an MRS Cluster.....	779
7.10.4 Changing the Passwords for System Users of an MRS Cluster.....	784
7.10.4.1 Changing or Resetting the Password for User admin of an MRS Cluster.....	784
7.10.4.2 Changing the Passwords for OS Users of an MRS Cluster Node.....	787
7.10.4.3 Changing the Password for the Kerberos Administrator of an MRS Cluster.....	788
7.10.4.4 Changing the Passwords for Manager Users of an MRS Cluster.....	790
7.10.4.5 Changing the Password for a Regular LDAP User of an MRS Cluster.....	792
7.10.4.6 Changing the LDAP Administrator Password for an MRS Cluster.....	794
7.10.4.7 Changing the Passwords for MRS Cluster Component Running Users.....	795
7.10.5 Changing the Passwords for Database Users of an MRS Cluster.....	798
7.10.5.1 Changing the Password for the OMS Database Administrator.....	798
7.10.5.2 Changing the Password for an OMS Database Access User.....	799
7.10.5.3 Changing the Passwords for Database Users of MRS Cluster Components.....	800
7.10.5.4 Resetting the MRS Component Database User Password.....	802
7.10.5.5 Resetting the Password for User omm in DBService.....	803
7.10.5.6 Changing the Password for User compdbuser of the DBService Database.....	804
7.11 Viewing and Configuring MRS Alarm Events.....	804
7.11.1 Viewing MRS Cluster Events.....	805

7.11.2 Viewing Alarms of an MRS Cluster.....	808
7.11.3 Configuring Alarm Thresholds for an MRS Cluster.....	813
7.11.4 Configuring Alarm Masking for an MRS Cluster.....	832
7.11.5 Connecting an MRS Cluster to SNMP to Report Alarms.....	833
7.11.6 Connecting an MRS Cluster to the Syslog Server to Report Alarms.....	835
7.11.7 Periodically Backing Up Alarm and Audit Information.....	840
7.11.8 Enabling the MRS Cluster Maintenance Mode to Disable Alarm Reporting.....	841
7.11.9 Configuring Notifications for MRS Cluster Alarms and Events.....	844
7.12 MRS Cluster Alarm Handling Reference.....	847
7.12.1 ALM-12001 Audit Log Dumping Failure.....	847
7.12.2 ALM-12004 Oldap Resource Abnormal.....	850
7.12.3 ALM-12005 OKerberos Resource Abnormal.....	851
7.12.4 ALM-12006 Node Fault.....	853
7.12.5 ALM-12007 Process Fault.....	857
7.12.6 ALM-12010 Manager Heartbeat Interruption Between the Active and Standby Nodes.....	861
7.12.7 ALM-12011 Manager Data Synchronization Exception Between the Active and Standby Nodes..	863
7.12.8 ALM-12012 NTP Service Is Abnormal.....	867
7.12.9 ALM-12014 Partition Lost.....	874
7.12.10 ALM-12015 Partition Filesystem Readonly.....	877
7.12.11 ALM-12016 CPU Usage Exceeds the Threshold.....	878
7.12.12 ALM-12017 Insufficient Disk Capacity.....	881
7.12.13 ALM-12018 Memory Usage Exceeds the Threshold.....	884
7.12.14 ALM-12027 Host PID Usage Exceeds the Threshold.....	886
7.12.15 ALM-12028 Number of Processes in the D State and Z State on a Host Exceeds the Threshold.	888
7.12.16 ALM-12033 Slow Disk Fault.....	890
7.12.17 ALM-12034 Periodical Backup Failure.....	897
7.12.18 ALM-12035 Unknown Data Status After Recovery Task Failure.....	900
7.12.19 ALM-12037 NTP Server Abnormal.....	902
7.12.20 ALM-12038 Monitoring Indicator Dumping Failure.....	905
7.12.21 ALM-12039 Active/Standby OMS Databases Not Synchronized.....	907
7.12.22 ALM-12040 Insufficient System Entropy.....	910
7.12.23 ALM-12041 Incorrect Permission on Key Files.....	913
7.12.24 ALM-12042 Incorrect Configuration of Key Files.....	915
7.12.25 ALM-12045 Read Packet Dropped Rate Exceeds the Threshold.....	917
7.12.26 ALM-12046 Write Packet Dropped Rate Exceeds the Threshold.....	921
7.12.27 ALM-12047 Read Packet Error Rate Exceeds the Threshold.....	924
7.12.28 ALM-12048 Write Packet Error Rate Exceeds the Threshold.....	926
7.12.29 ALM-12049 Network Read Throughput Rate Exceeds the Threshold.....	929
7.12.30 ALM-12050 Network Write Throughput Rate Exceeds the Threshold.....	932
7.12.31 ALM-12051 Disk Inode Usage Exceeds the Threshold.....	935
7.12.32 ALM-12052 TCP Temporary Port Usage Exceeds the Threshold.....	937
7.12.33 ALM-12053 Host File Handle Usage Exceeds the Threshold.....	940

7.12.34 ALM-12054 Invalid Certificate File.....	943
7.12.35 ALM-12055 Certificate File Is About to Expire.....	946
7.12.36 ALM-12057 Metadata Not Configured with the Task to Periodically Back Up Data to a Third-Party Server.....	949
7.12.37 ALM-12061 Process Usage Exceeds the Threshold.....	950
7.12.38 ALM-12062 OMS Parameter Configurations Mismatch with the Cluster Scale.....	954
7.12.39 ALM-12063 Unavailable Disk.....	956
7.12.40 ALM-12064 Host Random Port Range Conflicts with Cluster Used Port.....	958
7.12.41 ALM-12066 Trust Relationships Between Nodes Become Invalid.....	960
7.12.42 ALM-12067 Tomcat Resource Is Abnormal.....	964
7.12.43 ALM-12068 ACS Resource Exception.....	965
7.12.44 ALM-12069 AOS Resource Exception.....	967
7.12.45 ALM-12070 Controller Resource Is Abnormal.....	969
7.12.46 ALM-12071 Httpd Resource Is Abnormal.....	971
7.12.47 ALM-12072 FloatIP Resource Is Abnormal.....	973
7.12.48 ALM-12073 CEP Resource Is Abnormal.....	975
7.12.49 ALM-12074 FMS Resource Is Abnormal.....	977
7.12.50 ALM-12075 PMS Resource Is Abnormal.....	979
7.12.51 ALM-12076 GaussDB Resource Is Abnormal.....	981
7.12.52 ALM-12077 User omm Expired.....	983
7.12.53 ALM-12078 Password of User omm Expired.....	985
7.12.54 ALM-12079 User omm Is About to Expire.....	987
7.12.55 ALM-12080 Password of User omm Is About to Expire.....	988
7.12.56 ALM-12081 User ommdba Expired.....	990
7.12.57 ALM-12082 User ommdba Is About to Expire.....	992
7.12.58 ALM-12083 Password of User ommdba Is About to Expire.....	994
7.12.59 ALM-12084 Password of User ommdba Expired.....	995
7.12.60 ALM-12085 Service Audit Log Dump Failure.....	997
7.12.61 ALM-12087 System Is in the Upgrade Observation Period.....	1000
7.12.62 ALM-12089 Inter-Node Network Is Abnormal.....	1002
7.12.63 ALM-12091 Abnormal disaster Resources.....	1004
7.12.64 ALM-12099 core dump Occurred.....	1005
7.12.65 ALM-12100 AD Service Connection Failed.....	1007
7.12.66 ALM-12101 AZ Unhealthy.....	1010
7.12.67 ALM-12102 AZ HA Component Is Not Deployed Based on DR Requirements.....	1012
7.12.68 ALM-12103 Executor Resource Exception.....	1013
7.12.69 ALM-12104 Abnormal Knox Resources.....	1015
7.12.70 ALM-12110 Failed to get ECS temporary AK/SK.....	1016
7.12.71 ALM-12172 Failed to Report Metrics to Cloud Eye.....	1018
7.12.72 ALM-12180 Suspended Disk I/O.....	1020
7.12.73 ALM-12186 CGroup Task Usage Exceeds the Threshold.....	1024
7.12.74 ALM-12187 Failed to Expand Disk Partition Capacity.....	1026
7.12.75 ALM-12188 diskmgt Disk Monitoring Unavailable.....	1028

7.12.76 ALM-12190 Number of Knox Connections Exceeds the Threshold.....	1030
7.12.77 ALM-13000 ZooKeeper Service Unavailable.....	1032
7.12.78 ALM-13001 Available ZooKeeper Connections Are Insufficient.....	1035
7.12.79 ALM-13002 ZooKeeper Direct Memory Usage Exceeds the Threshold.....	1038
7.12.80 ALM-13003 GC Duration of the ZooKeeper Process Exceeds the Threshold.....	1040
7.12.81 ALM-13004 ZooKeeper Heap Memory Usage Exceeds the Threshold.....	1043
7.12.82 ALM-13005 Failed to Set the Quota of Top Directories of ZooKeeper Components.....	1046
7.12.83 ALM-13006 Znode Number or Capacity Exceeds the Threshold.....	1048
7.12.84 ALM-13007 Available ZooKeeper Client Connections Are Insufficient.....	1051
7.12.85 ALM-13008 ZooKeeper Znode Usage Exceeds the Threshold.....	1053
7.12.86 ALM-13009 ZooKeeper Znode Capacity Usage Exceeds the Threshold.....	1055
7.12.87 ALM-13010 Znode Usage of a Directory with Quota Configured Exceeds the Threshold.....	1057
7.12.88 ALM-14000 HDFS Service Unavailable.....	1059
7.12.89 ALM-14001 HDFS Disk Usage Exceeds the Threshold.....	1062
7.12.90 ALM-14002 DataNode Disk Usage Exceeds the Threshold.....	1065
7.12.91 ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold.....	1068
7.12.92 ALM-14006 Number of HDFS Files Exceeds the Threshold.....	1071
7.12.93 ALM-14007 NameNode Heap Memory Usage Exceeds the Threshold.....	1074
7.12.94 ALM-14008 DataNode Heap Memory Usage Exceeds the Threshold.....	1077
7.12.95 ALM-14009 Number of Dead DataNodes Exceeds the Threshold.....	1080
7.12.96 ALM-14010 NameService Service Is Abnormal.....	1083
7.12.97 ALM-14011 DataNode Data Directory Is Not Configured Properly.....	1087
7.12.98 ALM-14012 JournalNode Is Out of Synchronization.....	1091
7.12.99 ALM-14013 Failed to Update the NameNode FsImage File.....	1093
7.12.100 ALM-14014 NameNode GC Time Exceeds the Threshold.....	1099
7.12.101 ALM-14015 DataNode GC Time Exceeds the Threshold.....	1102
7.12.102 ALM-14016 DataNode Direct Memory Usage Exceeds the Threshold.....	1104
7.12.103 ALM-14017 NameNode Direct Memory Usage Exceeds the Threshold.....	1107
7.12.104 ALM-14018 NameNode Non-heap Memory Usage Exceeds the Threshold.....	1109
7.12.105 ALM-14019 DataNode Non-heap Memory Usage Exceeds the Threshold.....	1112
7.12.106 ALM-14020 Number of Entries in the HDFS Directory Exceeds the Threshold.....	1115
7.12.107 ALM-14021 NameNode Average RPC Processing Time Exceeds the Threshold.....	1118
7.12.108 ALM-14022 NameNode Average RPC Queuing Time Exceeds the Threshold.....	1123
7.12.109 ALM-14023 Percentage of Total Reserved Disk Space for Replicas Exceeds the Threshold.....	1128
7.12.110 ALM-14024 Tenant Space Usage Exceeds the Threshold.....	1130
7.12.111 ALM-14025 Tenant File Object Usage Exceeds the Threshold.....	1133
7.12.112 ALM-14026 Blocks on DataNode Exceed the Threshold.....	1135
7.12.113 ALM-14027 DataNode Disk Fault.....	1138
7.12.114 ALM-14028 Number of Blocks to Be Supplemented Exceeds the Threshold.....	1141
7.12.115 ALM-14029 Number of Blocks in a Replica Exceeds the Threshold.....	1143
7.12.116 ALM-14030 HDFS Allows Write of Single-Replica Data.....	1146
7.12.117 ALM-14031 DataNode Process Is Abnormal.....	1147

7.12.118 ALM-14032 JournalNode Process Is Abnormal.....	1149
7.12.119 ALM-14033 ZKFC Process Is Abnormal.....	1151
7.12.120 ALM-14034 Router Process Is Abnormal.....	1153
7.12.121 ALM-14035 HttpFS Process Is Abnormal.....	1155
7.12.122 ALM-16000 Percentage of Sessions Connected to the HiveServer to Maximum Number Allowed Exceeds the Threshold.....	1157
7.12.123 ALM-16001 Hive Warehouse Space Usage Exceeds the Threshold.....	1159
7.12.124 ALM-16002 Hive SQL Execution Success Rate Is Lower Than the Threshold.....	1161
7.12.125 ALM-16003 Background Thread Usage Exceeds the Threshold.....	1164
7.12.126 ALM-16004 Hive Service Unavailable.....	1167
7.12.127 ALM-16005 The Heap Memory Usage of the Hive Process Exceeds the Threshold.....	1171
7.12.128 ALM-16006 The Direct Memory Usage of the Hive Process Exceeds the Threshold.....	1175
7.12.129 ALM-16007 Hive GC Time Exceeds the Threshold.....	1179
7.12.130 ALM-16008 Non-Heap Memory Usage of the Hive Process Exceeds the Threshold.....	1183
7.12.131 ALM-16009 Map Number Exceeds the Threshold.....	1187
7.12.132 ALM-16045 Hive Data Warehouse Is Deleted.....	1189
7.12.133 ALM-16046 Hive Data Warehouse Permission Is Modified.....	1190
7.12.134 ALM-16047 HiveServer Has Been Deregistered from ZooKeeper.....	1192
7.12.135 ALM-16048 Tez or Spark Library Path Does Not Exist.....	1195
7.12.136 ALM-17003 Oozie Service Unavailable.....	1196
7.12.137 ALM-17004 Oozie Heap Memory Usage Exceeds the Threshold.....	1200
7.12.138 ALM-17005 Oozie Non Heap Memory Usage Exceeds the Threshold.....	1203
7.12.139 ALM-17006 Oozie Direct Memory Usage Exceeds the Threshold.....	1206
7.12.140 ALM-17007 Garbage Collection (GC) Time of the Oozie Process Exceeds the Threshold.....	1208
7.12.141 ALM-17008 Abnormal Connection Between Oozie and ZooKeeper.....	1211
7.12.142 ALM-17009 Abnormal Connection Between Oozie and DBService.....	1213
7.12.143 ALM-17010 Abnormal Connection Between Oozie and HDFS.....	1215
7.12.144 ALM-17011 Abnormal Connection Between Oozie and Yarn.....	1217
7.12.145 ALM-18000 Yarn Service Unavailable.....	1219
7.12.146 ALM-18002 NodeManager Heartbeat Lost.....	1221
7.12.147 ALM-18003 NodeManager Unhealthy.....	1224
7.12.148 ALM-18008 Heap Memory Usage of ResourceManager Exceeds the Threshold.....	1227
7.12.149 ALM-18009 Heap Memory Usage of JobHistoryServer Exceeds the Threshold.....	1231
7.12.150 ALM-18010 ResourceManager GC Time Exceeds the Threshold.....	1233
7.12.151 ALM-18011 NodeManager GC Time Exceeds the Threshold.....	1237
7.12.152 ALM-18012 JobHistoryServer GC Time Exceeds the Threshold.....	1240
7.12.153 ALM-18013 ResourceManager Direct Memory Usage Exceeds the Threshold.....	1242
7.12.154 ALM-18014 NodeManager Direct Memory Usage Exceeds the Threshold.....	1245
7.12.155 ALM-18015 JobHistoryServer Direct Memory Usage Exceeds the Threshold.....	1247
7.12.156 ALM-18016 Non Heap Memory Usage of ResourceManager Exceeds the Threshold.....	1250
7.12.157 ALM-18017 Non Heap Memory Usage of NodeManager Exceeds the Threshold.....	1253
7.12.158 ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold.....	1256
7.12.159 ALM-18019 Non Heap Memory Usage of JobHistoryServer Exceeds the Threshold.....	1258

7.12.160 ALM-18020 Yarn Task Execution Timeout.....	1261
7.12.161 ALM-18021 Mapreduce Service Unavailable.....	1263
7.12.162 ALM-18022 Insufficient Yarn Queue Resources.....	1266
7.12.163 ALM-18023 Number of Pending Yarn Tasks Exceeds the Threshold.....	1269
7.12.164 ALM-18024 Pending Yarn Memory Usage Exceeds the Threshold.....	1271
7.12.165 ALM-18025 Number of Terminated Yarn Tasks Exceeds the Threshold.....	1273
7.12.166 ALM-18026 Number of Failed Yarn Tasks Exceeds the Threshold.....	1275
7.12.167 ALM-19000 HBase Service Unavailable.....	1277
7.12.168 ALM-19006 HBase Replication Sync Failed.....	1283
7.12.169 ALM-19007 HBase GC Time Exceeds the Threshold.....	1287
7.12.170 ALM-19008 Heap Memory Usage of the HBase Process Exceeds the Threshold.....	1291
7.12.171 ALM-19009 Direct Memory Usage of the HBase Process Exceeds the Threshold.....	1294
7.12.172 ALM-19011 RegionServer Region Number Exceeds the Threshold.....	1298
7.12.173 ALM-19012 HBase System Table Directory or File Lost.....	1302
7.12.174 ALM-19013 Duration of Regions in transaction State Exceeds the Threshold.....	1304
7.12.175 ALM-19014 Capacity Quota Usage on ZooKeeper Exceeds the Threshold Severely.....	1307
7.12.176 ALM-19015 Quantity Quota Usage on ZooKeeper Exceeds the Threshold.....	1309
7.12.177 ALM-19016 Quantity Quota Usage on ZooKeeper Exceeds the Threshold Severely.....	1312
7.12.178 ALM-19017 Capacity Quota Usage on ZooKeeper Exceeds the Threshold.....	1315
7.12.179 ALM-19018 HBase Compaction Queue Size Exceeds the Threshold.....	1317
7.12.180 ALM-19019 Number of HBase HFiles to Be Synchronized Exceeds the Threshold.....	1319
7.12.181 ALM-19020 Number of HBase WAL Files to Be Synchronized Exceeds the Threshold.....	1322
7.12.182 ALM-19021 Handler Usage of RegionServer Exceeds the Threshold.....	1325
7.12.183 ALM-19022 HBase Hotspot Detection Is Unavailable.....	1328
7.12.184 ALM-19023 Region Traffic Restriction for HBase.....	1331
7.12.185 ALM-19024 RPC Requests P99 Latency on RegionServer Exceeds the Threshold.....	1333
7.12.186 ALM-19025 Damaged StoreFile in HBase.....	1336
7.12.187 ALM-19026 Damaged WAL Files in HBase.....	1338
7.12.188 ALM-20002 Hue Service Unavailable.....	1340
7.12.189 ALM-23001 Loader Service Unavailable.....	1342
7.12.190 ALM-23003 Loader Task Execution Failure.....	1346
7.12.191 ALM-23004 Loader Heap Memory Usage Exceeds the Threshold.....	1349
7.12.192 ALM-23005 Loader Non-Heap Memory Usage Exceeds the Threshold.....	1351
7.12.193 ALM-23006 Loader Direct Memory Usage Exceeds the Threshold.....	1354
7.12.194 ALM-23007 Garbage Collection (GC) Time of the Loader Process Exceeds the Threshold.....	1356
7.12.195 ALM-24000 Flume Service Unavailable.....	1359
7.12.196 ALM-24001 Flume Agent Exception.....	1360
7.12.197 ALM-24003 Flume Client Connection Interrupted.....	1364
7.12.198 ALM-24004 Exception Occurs When Flume Reads Data.....	1366
7.12.199 ALM-24005 Exception Occurs When Flume Transmits Data.....	1369
7.12.200 ALM-24006 Heap Memory Usage of Flume Server Exceeds the Threshold.....	1372
7.12.201 ALM-24007 Flume Server Direct Memory Usage Exceeds the Threshold.....	1375

7.12.202 ALM-24008 Flume Server Non Heap Memory Usage Exceeds the Threshold.....	1377
7.12.203 ALM-24009 Flume Server Garbage Collection (GC) Time Exceeds the Threshold.....	1380
7.12.204 ALM-24010 Flume Certificate File Is Invalid or Damaged.....	1383
7.12.205 ALM-24011 Flume Certificate File Is About to Expire.....	1385
7.12.206 ALM-24012 Flume Certificate File Has Expired.....	1387
7.12.207 ALM-24013 Flume MonitorServer Certificate File Is Invalid or Damaged.....	1390
7.12.208 ALM-24014 Flume MonitorServer Certificate Is About to Expire.....	1392
7.12.209 ALM-24015 Flume MonitorServer Certificate File Has Expired.....	1394
7.12.210 ALM-25000 LdapServer Service Unavailable.....	1397
7.12.211 ALM-25004 Abnormal LdapServer Data Synchronization.....	1399
7.12.212 ALM-25005 nscd Service Exception.....	1402
7.12.213 ALM-25006 Sssd Service Exception.....	1406
7.12.214 ALM-25007 Number of SlapdServer Connections Exceeds the Threshold.....	1409
7.12.215 ALM-25008 SlapdServer CPU Usage Exceeds the Threshold.....	1411
7.12.216 ALM-25500 KrbServer Service Unavailable.....	1414
7.12.217 ALM-26051 Storm Service Unavailable.....	1416
7.12.218 ALM-26052 Number of Available Supervisors of the Storm Service Is Less Than the Threshold	1418
7.12.219 ALM-26053 Storm Slot Usage Exceeds the Threshold.....	1420
7.12.220 ALM-26054 Nimbus Heap Memory Usage Exceeds the Threshold.....	1422
7.12.221 ALM-27001 DBService Service Unavailable.....	1425
7.12.222 ALM-27003 DBService Heartbeat Interruption Between the Active and Standby Nodes.....	1428
7.12.223 ALM-27004 Data Inconsistency Between Active and Standby DBServices.....	1430
7.12.224 ALM-27005 Database Connections Usage Exceeds the Threshold.....	1432
7.12.225 ALM-27006 Disk Space Usage of the Data Directory Exceeds the Threshold.....	1437
7.12.226 ALM-27007 Database Enters the Read-Only Mode.....	1439
7.12.227 ALM-29000 Impala Service Unavailable.....	1442
7.12.228 ALM-29004 Impalad Process Memory Usage Exceeds the Threshold.....	1445
7.12.229 ALM-29005 Number of JDBC Connections to Impalad Exceeds the Threshold.....	1447
7.12.230 ALM-29006 Number of ODBC Connections to Impalad Exceeds the Threshold.....	1449
7.12.231 ALM-29007 Impalad Process Memory Usage Exceeds the Threshold.....	1452
7.12.232 ALM-29008 Number of ODBC Connections to Impalad Exceeds the Threshold.....	1454
7.12.233 ALM-29010 Number of Queries Being Submitted by Impalad Exceeds the Threshold.....	1456
7.12.234 ALM-29011 Number of Queries Being Executed by Impalad Exceeds the Threshold.....	1458
7.12.235 ALM-29012 Number of Queries Being Waited by Impalad Exceeds the Threshold.....	1460
7.12.236 ALM-29013 Impalad FGC Time Exceeds the Threshold.....	1462
7.12.237 ALM-29014 Catalog FGC Time Exceeds the Threshold.....	1464
7.12.238 ALM-29015 Catalog Process Memory Usage Exceeds the Threshold.....	1466
7.12.239 ALM-29016 Impalad Instance in the Sub-healthy State.....	1468
7.12.240 ALM-29100 Kudu Service Unavailable.....	1470
7.12.241 ALM-29104 Tserver Process Memory Usage Exceeds the Threshold.....	1471
7.12.242 ALM-29106 Tserver Process CPU Usage Exceeds the Threshold.....	1473
7.12.243 ALM-29107 Tserver Process Memory Usage Exceeds the Threshold.....	1475

7.12.244 ALM-38000 Kafka Service Unavailable.....	1476
7.12.245 ALM-38001 Insufficient Kafka Disk Capacity.....	1479
7.12.246 ALM-38002 Kafka Heap Memory Usage Exceeds the Threshold.....	1484
7.12.247 ALM-38004 Kafka Direct Memory Usage Exceeds the Threshold.....	1487
7.12.248 ALM-38005 GC Duration of the Broker Process Exceeds the Threshold.....	1490
7.12.249 ALM-38006 Percentage of Kafka Partitions That Are Not Completely Synchronized Exceeds the Threshold.....	1493
7.12.250 ALM-38007 Status of Kafka Default User Is Abnormal.....	1495
7.12.251 ALM-38008 Abnormal Kafka Data Directory Status.....	1497
7.12.252 ALM-38009 Busy Broker Disk I/Os (Applicable to Versions Later Than MRS 3.1.0).....	1499
7.12.253 ALM-38009 Kafka Topic Overload (Applicable to MRS 3.1.0 and Earlier Versions).....	1502
7.12.254 ALM-38010 Topics with Single Replica.....	1505
7.12.255 ALM-38011 User Connection Usage on Broker Exceeds the Threshold.....	1507
7.12.256 ALM-43001 Spark2x Service Unavailable.....	1511
7.12.257 ALM-43006 Heap Memory Usage of the JobHistory2x Process Exceeds the Threshold.....	1514
7.12.258 ALM-43007 Non-Heap Memory Usage of the JobHistory2x Process Exceeds the Threshold....	1517
7.12.259 ALM-43008 The Direct Memory Usage of the JobHistory2x Process Exceeds the Threshold....	1520
7.12.260 ALM-43009 JobHistory2x Process GC Time Exceeds the Threshold.....	1524
7.12.261 ALM-43010 Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold.....	1527
7.12.262 ALM-43011 Non-Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold... 1531	
7.12.263 ALM-43012 Direct Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold 1534	
7.12.264 ALM-43013 JDBCServer2x Process GC Time Exceeds the Threshold.....	1538
7.12.265 ALM-43017 JDBCServer2x Process Full GC Number Exceeds the Threshold.....	1540
7.12.266 ALM-43018 JobHistory2x Process Full GC Number Exceeds the Threshold.....	1543
7.12.267 ALM-43019 Heap Memory Usage of the IndexServer2x Process Exceeds the Threshold.....	1546
7.12.268 ALM-43020 Non-Heap Memory Usage of the IndexServer2x Process Exceeds the Threshold..	1550
7.12.269 ALM-43021 Direct Memory Usage of the IndexServer2x Process Exceeds the Threshold.....	1554
7.12.270 ALM-43022 IndexServer2x Process GC Time Exceeds the Threshold.....	1557
7.12.271 ALM-43023 IndexServer2x Process Full GC Number Exceeds the Threshold.....	1560
7.12.272 ALM-44000 Presto Service Unavailable.....	1563
7.12.273 ALM-44004 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold.....	1564
7.12.274 ALM-44005 Presto Coordinator Process GC Time Exceeds the Threshold.....	1566
7.12.275 ALM-44006 Presto Worker Process GC Time Exceeds the Threshold.....	1567
7.12.276 ALM-45000 HetuEngine Service Unavailable.....	1569
7.12.277 ALM-45001 Faulty HetuEngine Compute Instances.....	1573
7.12.278 ALM-45003 HetuEngine QAS Disk Capacity Is Insufficient.....	1575
7.12.279 ALM-45175 Average Time for Calling OBS Metadata APIs Is Greater than the Threshold.....	1578
7.12.280 ALM-45176 Success Rate of Calling OBS Metadata APIs Is Lower than the Threshold.....	1580
7.12.281 ALM-45177 Success Rate of Calling OBS Data Read APIs Is Lower than the Threshold.....	1583
7.12.282 ALM-45178 Success Rate of Calling OBS Data Write APIs Is Lower Than the Threshold.....	1585
7.12.283 ALM-45179 Number of Failed OBS readFully API Calls Exceeds the Threshold.....	1588
7.12.284 ALM-45180 Number of Failed OBS read API Calls Exceeds the Threshold.....	1590
7.12.285 ALM-45181 Number of Failed OBS write API Calls Exceeds the Threshold.....	1592

7.12.286 ALM-45182 Number of Throttled OBS Operations Exceeds the Threshold.....	1594
7.12.287 ALM-45275 Ranger Service Unavailable.....	1596
7.12.288 ALM-45276 Abnormal RangerAdmin Status.....	1598
7.12.289 ALM-45277 RangerAdmin Heap Memory Usage Exceeds the Threshold.....	1600
7.12.290 ALM-45278 RangerAdmin Direct Memory Usage Exceeds the Threshold.....	1602
7.12.291 ALM-45279 RangerAdmin Non Heap Memory Usage Exceeds the Threshold.....	1605
7.12.292 ALM-45280 RangerAdmin GC Duration Exceeds the Threshold.....	1608
7.12.293 ALM-45281 UserSync Heap Memory Usage Exceeds the Threshold.....	1610
7.12.294 ALM-45282 UserSync Direct Memory Usage Exceeds the Threshold.....	1613
7.12.295 ALM-45283 UserSync Non Heap Memory Usage Exceeds the Threshold.....	1616
7.12.296 ALM-45284 UserSync Garbage Collection (GC) Time Exceeds the Threshold.....	1618
7.12.297 ALM-45285 TagSync Heap Memory Usage Exceeds the Threshold.....	1621
7.12.298 ALM-45286 TagSync Direct Memory Usage Exceeds the Threshold.....	1624
7.12.299 ALM-45287 TagSync Non Heap Memory Usage Exceeds the Threshold.....	1626
7.12.300 ALM-45288 TagSync Garbage Collection (GC) Time Exceeds the Threshold.....	1629
7.12.301 ALM-45289 PolicySync Heap Memory Usage Exceeds the Threshold.....	1631
7.12.302 ALM-45290 PolicySync Direct Memory Usage Exceeds the Threshold.....	1634
7.12.303 ALM-45291 PolicySync Non-Heap Memory Usage Exceeds the Threshold.....	1636
7.12.304 ALM-45292 PolicySync GC Duration Exceeds the Threshold.....	1639
7.12.305 ALM-45325 Presto Service Unavailable.....	1641
7.12.306 ALM-45326 Number of Presto Coordinator Threads Exceeds the Threshold.....	1643
7.12.307 ALM-45327 Presto Coordinator Process GC Time Exceeds the Threshold.....	1644
7.12.308 ALM-45328 Presto Worker Process GC Time Exceeds the Threshold.....	1646
7.12.309 ALM-45329 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold.....	1648
7.12.310 ALM-45330 Number of Presto Worker Threads Exceeds the Threshold.....	1650
7.12.311 ALM-45331 Number of Presto Worker1 Threads Exceeds the Threshold.....	1651
7.12.312 ALM-45332 Number of Presto Worker2 Threads Exceeds the Threshold.....	1653
7.12.313 ALM-45333 Number of Presto Worker3 Threads Exceeds the Threshold.....	1655
7.12.314 ALM-45334 Number of Presto Worker4 Threads Exceeds the Threshold.....	1657
7.12.315 ALM-45335 Presto Worker1 Process GC Time Exceeds the Threshold.....	1658
7.12.316 ALM-45336 Presto Worker2 Process GC Time Exceeds the Threshold.....	1660
7.12.317 ALM-45337 Presto Worker3 Process GC Time Exceeds the Threshold.....	1662
7.12.318 ALM-45338 Presto Worker4 Process GC Time Exceeds the Threshold.....	1664
7.12.319 ALM-45425 ClickHouse Service Unavailable.....	1666
7.12.320 ALM-45426 ClickHouse Service Quantity Quota Usage in ZooKeeper Exceeds the Threshold..	1669
7.12.321 ALM-45427 ClickHouse Service Capacity Quota Usage in ZooKeeper Exceeds the Threshold..	1672
7.12.322 ALM-45428 ClickHouse Disk I/O Exception.....	1674
7.12.323 ALM-45429 Table Metadata Synchronization Failed on the Added ClickHouse Node.....	1677
7.12.324 ALM-45430 Permission Metadata Synchronization Failed on the Added ClickHouse Node.....	1679
7.12.325 ALM-45431 Improper ClickHouse Instance Distribution for Topology Allocation.....	1681
7.12.326 ALM-45432 ClickHouse User Synchronization Process Fails.....	1683
7.12.327 ALM-45433 ClickHouse AZ Topology Exception.....	1685

7.12.328 ALM-45434 A Single Replica Exists in the ClickHouse Data Table.....	1688
7.12.329 ALM-45435 Inconsistent Metadata of ClickHouse Tables.....	1690
7.12.330 ALM-45436 Skew ClickHouse Table Data.....	1692
7.12.331 ALM-45437 Excessive Parts in the ClickHouse Table.....	1694
7.12.332 ALM-45438 ClickHouse Disk Usage Exceeds 80%.....	1696
7.12.333 ALM-45439 ClickHouse Node Enters the Read-Only Mode.....	1698
7.12.334 ALM-45440 Inconsistency Between ClickHouse Replicas.....	1699
7.12.335 ALM-45441 Zookeeper Disconnected.....	1703
7.12.336 ALM-45442 Too Many Concurrent SQL Statements.....	1705
7.12.337 ALM-45443 Slow SQL Queries in the Cluster.....	1707
7.12.338 ALM-45444 Abnormal ClickHouse Process.....	1709
7.12.339 ALM-45585 IoTDB Service Unavailable.....	1711
7.12.340 ALM-45586 IoTDBServer Heap Memory Usage Exceeds the Threshold.....	1713
7.12.341 ALM-45587 IoTDBServer GC Duration Exceeds the Threshold.....	1715
7.12.342 ALM-45588 IoTDBServer Direct Memory Usage Exceeds the Threshold.....	1717
7.12.343 ALM-45589 ConfigNode Heap Memory Usage Exceeds the Threshold.....	1719
7.12.344 ALM-45590 ConfigNode GC Duration Exceeds the Threshold.....	1722
7.12.345 ALM-45591 ConfigNode Direct Memory Usage Exceeds the Threshold.....	1724
7.12.346 ALM-45592 IoTDBServer RPC Execution Duration Exceeds the Threshold.....	1726
7.12.347 ALM-45593 IoTDBServer Flush Execution Duration Exceeds the Threshold.....	1728
7.12.348 ALM-45594 IoTDBServer Intra-Space Merge Duration Exceeds the Threshold.....	1729
7.12.349 ALM-45595 IoTDBServer Cross-Space Merge Duration Exceeds the Threshold.....	1731
7.12.350 ALM-45596 Procedure Execution Failed.....	1732
7.12.351 ALM-45615 CDL Service Unavailable.....	1734
7.12.352 ALM-45616 CDL Job Execution Exception.....	1736
7.12.353 ALM-45617 Data Queued in the CDL Replication Slot Exceeds the Threshold.....	1738
7.12.354 ALM-45635 FlinkServer Job Execution Failure.....	1740
7.12.355 ALM-45636 Flink Job Checkpoints Keep Failing.....	1743
7.12.356 ALM-45636 Number of Consecutive Checkpoint Failures of a Flink Job Exceeds the Threshold.....	1745
7.12.357 ALM-45637 FlinkServer Task Is Continuously Under Back Pressure.....	1748
7.12.358 ALM-45638 Number of Restarts After FlinkServer Job Failures Exceeds the Threshold.....	1750
7.12.359 ALM-45638 Number of Restarts After Flink Job Failures Exceeds the Threshold.....	1753
7.12.360 ALM-45639 Checkpointing of a Flink Job Times Out.....	1755
7.12.361 ALM-45640 FlinkServer Heartbeat Interruption Between the Active and Standby Nodes.....	1758
7.12.362 ALM-45641 Data Synchronization Exception Between the Active and Standby FlinkServer Nodes	1760
7.12.363 ALM-45642 RocksDB Continuously Triggers Write Traffic Limiting.....	1764
7.12.364 ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold.....	1768
7.12.365 ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold.....	1771
7.12.366 ALM-45645 Pending Flush Size of RocksDB Continuously Exceeds the Threshold.....	1775
7.12.367 ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold.....	1778
7.12.368 ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold	1782

7.12.369 ALM-45648 RocksDB Frequently Encounters Write-Stopped.....	1785
7.12.370 ALM-45649 P95 Latency of RocksDB Get Requests Continuously Exceeds the Threshold.....	1789
7.12.371 ALM-45650 P95 Latency of RocksDB Write Requests Continuously Exceeds the Threshold.....	1793
7.12.372 ALM-45652 Flink Service Unavailable.....	1797
7.12.373 ALM-45653 Invalid Flink HA Certificate File.....	1799
7.12.374 ALM-45654 Flink HA Certificate Is About to Expire.....	1801
7.12.375 ALM-45655 Flink HA Certificate File Has Expired.....	1803
7.12.376 ALM-45736 Guardian Service Unavailable.....	1805
7.12.377 ALM-45737 TokenServer Heap Memory Usage Exceeds the Threshold.....	1807
7.12.378 ALM-45738 TokenServer Direct Memory Usage Exceeds the Threshold.....	1810
7.12.379 ALM-45739 TokenServer Non-Heap Memory Usage Exceeds the Threshold.....	1813
7.12.380 ALM-45740 TokenServer GC Duration Exceeds the Threshold.....	1815
7.12.381 ALM-45741 Failed to Call the ECS securitykey API.....	1818
7.12.382 ALM-45742 Failed to Call the ECS Metadata API.....	1820
7.12.383 ALM-45743 Failed to Call the IAM API.....	1821
7.12.384 ALM-50201 Doris Service Unavailable.....	1822
7.12.385 ALM-50202 FE CPU Usage Exceeds the Threshold.....	1824
7.12.386 ALM-50203 FE Memory Usage Exceeds the Threshold.....	1826
7.12.387 ALM-50205 BE CPU Usage Exceeds the Threshold.....	1827
7.12.388 ALM-50206 BE Memory Usage Exceeds the Threshold.....	1829
7.12.389 ALM-50207 Ratio of Connections to the FE MySQL Port to the Maximum Connections Allowed Exceeds the Threshold.....	1831
7.12.390 ALM-50208 Failures to Clear Historical Metadata Image Files Exceed the Threshold.....	1833
7.12.391 ALM-50209 Failures to Generate Metadata Image Files Exceed the Threshold.....	1835
7.12.392 ALM-50210 Maximum Compaction Score of All BE Nodes Exceeds the Threshold.....	1836
7.12.393 ALM-50211 FE Queue Length of BE Periodic Report Tasks Exceeds the Threshold.....	1838
7.12.394 ALM-50212 Accumulated Old-Generation GC Duration of the FE Process Exceeds the Threshold	1840
7.12.395 ALM-50213 Number of Tasks Queuing in the FE Thread Pool for Interacting with BE Exceeds the Threshold.....	1842
7.12.396 ALM-50214 Number of Tasks Queuing in the FE Thread Pool for Task Processing Exceeds the Threshold.....	1844
7.12.397 ALM-50215 Longest Duration of RPC Requests Received by Each FE Thrift Method Exceeds the Threshold.....	1846
7.12.398 ALM-50216 Memory Usage of the FE Node Exceeds the Threshold.....	1848
7.12.399 ALM-50217 Heap Memory Usage of the FE Node Exceeds the Threshold.....	1850
7.12.400 ALM-50219 Length of the Queue in the Thread Pool for Query Execution Exceeds the Threshold	1853
7.12.401 ALM-50220 Error Rate of TCP Packet Receiving Exceeds the Threshold.....	1855
7.12.402 ALM-50221 BE Data Disk Usage Exceeds the Threshold.....	1856
7.12.403 ALM-50222 Disk Status of a Specified Data Directory on BE Is Abnormal.....	1858
7.12.404 ALM-50223 Maximum Memory Required by BE Is Greater Than the Remaining Memory of the Machine.....	1860
7.12.405 ALM-50224 Failures a Certain Task Type on BE Are Increasing.....	1862

7.12.406 ALM-50225 FE Instance Fault.....	1864
7.12.407 ALM-50226 BE Instance Fault.....	1865
7.12.408 ALM-50401 Number of JobServer Jobs Waiting to Be Executed Exceeds the Threshold.....	1867
7.12.409 ALM-50402 JobGateway Service Unavailable.....	1869
7.12.410 ALM-12001 Audit Log Dump Failure (For MRS 2.x or Earlier).....	1871
7.12.411 ALM-12002 HA Resource Abnormal (For MRS 2.x or Earlier).....	1872
7.12.412 ALM-12004 Oldap Resource Abnormal (For MRS 2.x or Earlier).....	1875
7.12.413 ALM-12005 OKerberos Resource Abnormal (For MRS 2.x or Earlier).....	1876
7.12.414 ALM-12006 Node Fault (For MRS 2.x or Earlier).....	1878
7.12.415 ALM-12007 Process Fault (For MRS 2.x or Earlier).....	1879
7.12.416 ALM-12010 Manager Heartbeat Interruption Between the Active and Standby Nodes (For MRS 2.x or Earlier).....	1881
7.12.417 ALM-12011 Data Synchronization Exception Between the Active and Standby Manager Nodes (For MRS 2.x or Earlier).....	1883
7.12.418 ALM-12012 NTP Service Abnormal (For MRS 2.x or Earlier).....	1884
7.12.419 ALM-12014 Device Partition Lost (For MRS 2.x or Earlier).....	1887
7.12.420 ALM-12015 Device Partition File System Read-Only (For MRS 2.x or Earlier).....	1889
7.12.421 ALM-12016 CPU Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1891
7.12.422 ALM-12017 Insufficient Disk Capacity (For MRS 2.x or Earlier).....	1892
7.12.423 ALM-12018 Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1895
7.12.424 ALM-12027 Host PID Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1896
7.12.425 ALM-12028 Number of Processes in the D State on the Host Exceeds the Threshold (For MRS 2.x or Earlier).....	1898
7.12.426 ALM-12031 User omm or Password Is About to Expire (For MRS 2.x or Earlier).....	1900
7.12.427 ALM-12032 User ommdba or Password Is About to Expire (For MRS 2.x or Earlier).....	1901
7.12.428 ALM-12033 Slow Disk Fault (For MRS 2.x or Earlier).....	1903
7.12.429 ALM-12034 Periodic Backup Failure (For MRS 2.x or Earlier).....	1910
7.12.430 ALM-12035 Unknown Data Status After Recovery Task Failure (For MRS 2.x or Earlier).....	1912
7.12.431 ALM-12037 NTP Server Abnormal (For MRS 2.x or Earlier).....	1913
7.12.432 ALM-12038 Monitoring Indicator Dump Failure (For MRS 2.x or Earlier).....	1915
7.12.433 ALM-12039 GaussDB Data Is Not Synchronized (For MRS 2.x or Earlier).....	1917
7.12.434 ALM-12040 Insufficient System Entropy (For MRS 2.x or Earlier).....	1920
7.12.435 ALM-12041 Permission of Key Files Is Abnormal (For MRS 2.x or Earlier).....	1921
7.12.436 ALM-12042 Key File Configurations Are Abnormal (For MRS 2.x or Earlier).....	1923
7.12.437 ALM-12043 DNS Parsing Duration Exceeds the Threshold (For MRS 2.x or Earlier).....	1924
7.12.438 ALM-12045 Read Packet Dropped Rate Exceeds the Threshold (For MRS 2.x or Earlier).....	1927
7.12.439 ALM-12046 Write Packet Dropped Rate Exceeds the Threshold (For MRS 2.x or Earlier).....	1931
7.12.440 ALM-12047 Read Packet Error Rate Exceeds the Threshold (For MRS 2.x or Earlier).....	1933
7.12.441 ALM-12048 Write Packet Error Rate Exceeds the Threshold (For MRS 2.x or Earlier).....	1935
7.12.442 ALM-12049 Read Throughput Rate Exceeds the Threshold (For MRS 2.x or Earlier).....	1937
7.12.443 ALM-12050 Write Throughput Rate Exceeds the Threshold (For MRS 2.x or Earlier).....	1939
7.12.444 ALM-12051 Disk Inode Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1941
7.12.445 ALM-12052 Usage of Temporary TCP Ports Exceeds the Threshold (For MRS 2.x or Earlier)....	1942

7.12.446 ALM-12053 File Handle Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1945
7.12.447 ALM-12054 Invalid Certificate File (For MRS 2.x or Earlier).....	1946
7.12.448 ALM-12055 Certificate File Is About to Expire (For MRS 2.x or Earlier).....	1949
7.12.449 ALM-12180 Disk Card I/O (For MRS 2.x or Earlier).....	1951
7.12.450 ALM-12357 Failed to Export Audit Logs to OBS (For MRS 2.x or Earlier).....	1955
7.12.451 ALM-13000 ZooKeeper Service Unavailable (For MRS 2.x or Earlier).....	1956
7.12.452 ALM-13001 Available ZooKeeper Connections Are Insufficient (For MRS 2.x or Earlier).....	1959
7.12.453 ALM-13002 ZooKeeper Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1961
7.12.454 ALM-14000 HDFS Service Unavailable (For MRS 2.x or Earlier).....	1963
7.12.455 ALM-14001 HDFS Disk Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1965
7.12.456 ALM-14002 DataNode Disk Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1967
7.12.457 ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold (For MRS 2.x or Earlier).....	1968
7.12.458 ALM-14004 Number of Damaged HDFS Blocks Exceeds the Threshold (For MRS 2.x or Earlier)	1970
7.12.459 ALM-14006 Number of HDFS Files Exceeds the Threshold (For MRS 2.x or Earlier).....	1971
7.12.460 ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier)	1973
7.12.461 ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier)	1974
7.12.462 ALM-14009 Number of Faulty DataNodes Exceeds the Threshold (For MRS 2.x or Earlier).....	1976
7.12.463 ALM-14010 NameService Is Abnormal (For MRS 2.x or Earlier).....	1978
7.12.464 ALM-14011 HDFS DataNode Data Directory Is Not Configured Properly (For MRS 2.x or Earlier)	1981
7.12.465 ALM-14012 HDFS Journalnode Data Is Not Synchronized (For MRS 2.x or Earlier).....	1984
7.12.466 ALM-16000 Percentage of Sessions Connected to the HiveServer to the Maximum Number Allowed Exceeds the Threshold (For MRS 2.x or Earlier).....	1986
7.12.467 ALM-16001 Hive Warehouse Space Usage Exceeds the Threshold (For MRS 2.x or Earlier).....	1987
7.12.468 ALM-16002 Hive SQL Execution Success Rate Is Lower Than the Threshold (For MRS 2.x or Earlier).....	1989
7.12.469 ALM-16004 Hive Service Unavailable (For MRS 2.x or Earlier).....	1992
7.12.470 ALM-16005 Number of Failed Hive SQL Executions in the Last Period Exceeds the Threshold (For MRS 2.x or Earlier).....	1995
7.12.471 ALM-18000 Yarn Service Unavailable (For MRS 2.x or Earlier).....	1996
7.12.472 ALM-18002 NodeManager Heartbeat Lost (For MRS 2.x or Earlier).....	1998
7.12.473 ALM-18003 NodeManager Unhealthy (For MRS 2.x or Earlier).....	2000
7.12.474 ALM-18004 NodeManager Disk Usability Ratio Is Lower Than the Threshold (For MRS 2.x or Earlier).....	2001
7.12.475 ALM-18006 MapReduce Job Execution Timeout (For MRS 2.x or Earlier).....	2002
7.12.476 ALM-18008 Heap Memory Usage of Yarn ResourceManager Exceeds the Threshold (For MRS 2.x or Earlier).....	2004
7.12.477 ALM-18009 Heap Memory Usage of MapReduce JobHistoryServer Exceeds the Threshold (For MRS 2.x or Earlier).....	2006
7.12.478 ALM-18010 Number of Pending Yarn Tasks Exceeds the Threshold (For MRS 2.x or Earlier)...	2007
7.12.479 ALM-18011 Memory of Pending Yarn Tasks Exceeds the Threshold (For MRS 2.x or Earlier)...	2009
7.12.480 ALM-18012 Number of Terminated Yarn Tasks in the Last Period Exceeds the Threshold (For MRS 2.x or Earlier).....	2011

7.12.481 ALM-18013 Number of Failed Yarn Tasks in the Last Period Exceeds the Threshold (For MRS 2.x or Earlier).....	2012
7.12.482 ALM-19000 HBase Service Unavailable (For MRS 2.x or Earlier).....	2013
7.12.483 ALM-19006 HBase Replication Sync Failed (For MRS 2.x or Earlier).....	2014
7.12.484 ALM-19007 HBase Merge Queue Exceeds the Threshold (for 2.x and Earlier Versions).....	2017
7.12.485 ALM-20002 Hue Service Unavailable (For MRS 2.x or Earlier).....	2018
7.12.486 ALM-23001 Loader Service Unavailable (For MRS 2.x or Earlier).....	2021
7.12.487 ALM-24000 Flume Service Unavailable (For MRS 2.x or Earlier).....	2024
7.12.488 ALM-24001 Flume Agent Is Abnormal (For MRS 2.x or Earlier).....	2026
7.12.489 ALM-24003 Flume Client Connection Interrupted (For MRS 2.x or Earlier).....	2028
7.12.490 ALM-24004 Flume Fails to Read Data (For MRS 2.x or Earlier).....	2030
7.12.491 ALM-24005 Data Transmission by Flume Is Abnormal (For MRS 2.x or Earlier).....	2032
7.12.492 ALM-25000 LdapServer Service Unavailable (For MRS 2.x or Earlier).....	2034
7.12.493 ALM-25004 Abnormal LdapServer Data Synchronization (For MRS 2.x or Earlier).....	2036
7.12.494 ALM-25500 KrbServer Service Unavailable (For MRS 2.x or Earlier).....	2038
7.12.495 ALM-26051 Storm Service Unavailable (For MRS 2.x or Earlier).....	2040
7.12.496 ALM-26052 Number of Available Supervisors in Storm Is Lower Than the Threshold (For MRS 2.x or Earlier).....	2042
7.12.497 ALM-26053 Slot Usage of Storm Exceeds the Threshold (For MRS 2.x or Earlier).....	2044
7.12.498 ALM-26054 Heap Memory Usage of Storm Nimbus Exceeds the Threshold (For MRS 2.x or Earlier).....	2046
7.12.499 ALM-27001 DBService Unavailable (For MRS 2.x or Earlier).....	2047
7.12.500 ALM-27003 DBService Heartbeat Interruption Between the Active and Standby Nodes (For MRS 2.x or Earlier).....	2050
7.12.501 ALM-27004 Data Inconsistency Between Active and Standby DBServices (For MRS 2.x or Earlier).....	2051
7.12.502 ALM-28001 Spark Service Unavailable (For MRS 2.x or Earlier).....	2054
7.12.503 ALM-38000 Kafka Service Unavailable (For MRS 2.x or Earlier).....	2055
7.12.504 ALM-38001 Insufficient Kafka Disk Capacity (For MRS 2.x or Earlier).....	2057
7.12.505 ALM-38002 Heap Memory Usage of Kafka Exceeds the Threshold (For MRS 2.x or Earlier)....	2060
7.12.506 ALM-43001 Spark Service Unavailable (For MRS 2.x or Earlier).....	2062
7.12.507 ALM-43006 Heap Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.x or Earlier).....	2063
7.12.508 ALM-43007 Non-Heap Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.x or Earlier).....	2065
7.12.509 ALM-43008 Direct Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.x or Earlier).....	2066
7.12.510 ALM-43009 JobHistory GC Time Exceeds the Threshold (For MRS 2.x or Earlier).....	2068
7.12.511 ALM-43010 Heap Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.x or Earlier).....	2069
7.12.512 ALM-43011 Non-Heap Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.x or Earlier).....	2071
7.12.513 ALM-43012 Direct Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.x or Earlier).....	2072
7.12.514 ALM-43013 JDBCServer GC Time Exceeds the Threshold (For MRS 2.x or Earlier).....	2074

7.12.515 ALM-44004 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold (For MRS 2.x or Earlier).....	2075
7.12.516 ALM-44005 Presto Coordinator Process GC Time Exceeds the Threshold (For MRS 2.x or Earlier)	2077
7.12.517 ALM-44006 Presto Worker Process GC Time Exceeds the Threshold (For MRS 2.x or Earlier)..	2078
7.12.518 ALM-45325 Presto Service Unavailable (For MRS 2.x or Earlier).....	2080
7.13 Configuring Remote O&M for an MRS Cluster.....	2081
7.14 Common Ports for MRS Cluster Services.....	2082
8 Configuring Storage-Compute Decoupling for an MRS Cluster.....	2102
8.1 Configuration Process.....	2102
8.2 Interconnecting an MRS Cluster with OBS Using an IAM Agency.....	2104
8.2.1 Interconnecting an MRS Cluster with OBS Using an IAM Agency.....	2104
8.2.2 Configuring the Policy for Clearing Recycle Bin Directories of MRS Cluster Components.....	2109
8.2.3 Example for Interconnecting a Cluster Service with OBS.....	2112
8.2.3.1 Interconnecting Flink with OBS Using an IAM Agency.....	2112
8.2.3.2 Interconnecting Flume with OBS Using an IAM Agency.....	2113
8.2.3.3 Interconnecting HDFS with OBS Using an IAM Agency.....	2114
8.2.3.4 Interconnecting Hive with OBS Using an IAM Agency.....	2115
8.2.3.5 Interconnecting Hudi with OBS Using an IAM Agency.....	2119
8.2.3.6 Interconnecting MapReduce with OBS Using an IAM Agency.....	2121
8.2.3.7 Interconnecting Presto with OBS Using an IAM Agency.....	2122
8.2.3.8 Interconnecting Spark with OBS Using an IAM Agency.....	2124
8.2.3.9 Interconnecting Sqoop with OBS Using an IAM Agency.....	2127
8.2.4 Configuring Fine-Grained OBS Access Permissions for MRS Cluster Users.....	2132
8.3 FAQ About Decoupled Storage and Compute.....	2139
8.3.1 How Do I Read Encrypted OBS Data When Running an MRS Job?.....	2139
8.3.2 Example Application Development for Interconnecting HDFS with OBS.....	2147
8.3.3 How Do I Connect an MRS Cluster Client to OBS Using an AK/SK Pair?.....	2149
8.3.4 How Do I Access OBS Using an MRS Client Installed Outside a Cluster?.....	2155
8.3.5 Accessing an MRS Cluster's Manager (Version 2.x or Earlier).....	2157
8.3.6 How Do I Handle Abnormal Status of Core Nodes in an MRS Cluster After Successful Expansion?.....	2163

1 Preparations

1.1 Configuring MRS Cloud Service Authorization

You can interact with MRS clusters on the MRS console, monitor their status, and perform management operations. However, you need to obtain service authorization before using MRS for the first time.

After permission assignment, MRS creates an agency named **mrs_admin_agency** in Identity and Access Management (IAM). After the agency is created, do not modify or delete it. Deleting the agency or the Tenant Administrator role in the agency will automatically cancel the permission assignment. If permission assignment is canceled, functions such as cluster creation, cluster scale-in/out, Master node specification upgrade, auto scaling, cluster name modification, and IAM user synchronization will be affected, and the cluster running status cannot be monitored.

Registering with Huawei Cloud

If you have registered with Huawei Cloud, log in to the management console and access your MRS. If you do not have an account, register one with Huawei Cloud. After the registration, your account can be used to access all public cloud services, including your MRS.

Step 1 Visit [Huawei Cloud](#).

Step 2 Click **Register** and complete the registration as instructed.

After you have registered, the system automatically redirects you to your personal information page.

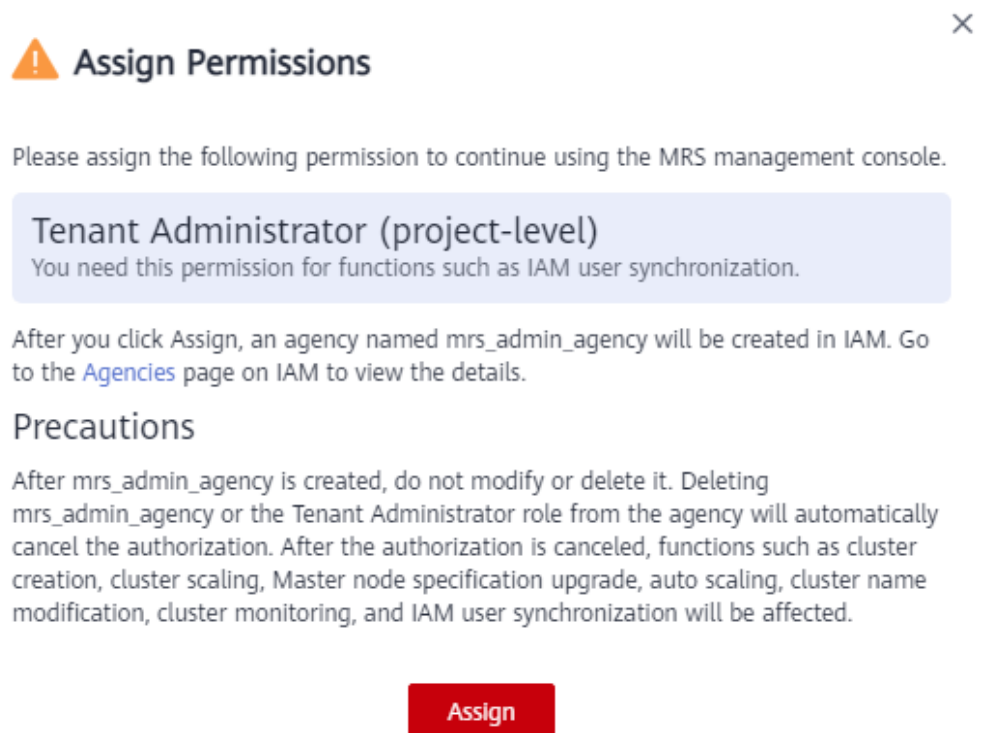
----End

Procedure

Step 1 Log in to the management console.

Step 2 Choose **Analytics > MapReduce Service** in the service list. The **Assign Permissions** page is displayed.

Figure 1-1 Assigning permissions



Step 3 Click **Assign**.

After assignment is agreed, an agency named **`mrs_admin_agency`** will be created in IAM. Do not modify or delete the agency after it is created. After the agency is created, you can use MRS.

 **NOTE**

After assignment, if the agency fails to be created, it is probably because the number of agencies already reaches the upper limit. In this case, log in to the IAM console and delete unnecessary agencies, or contact the administrator to increase the agency quota.

----End

1.2 Creating an IAM User and Granting MRS Permissions

Use **IAM** to implement fine-grained permission control over your MRS. With IAM, you can:

- Create IAM users under your Huawei Cloud account for employees based on your enterprise's organizational structure so that each employee is allowed to access MRS resources using their unique security credential (IAM user).
- Grant only the permissions required for users to perform a specific task.
- Entrust a Huawei Cloud account or cloud service to perform efficient O&M on your MRS resources.

If your Huawei Cloud account does not require IAM users, skip this section.

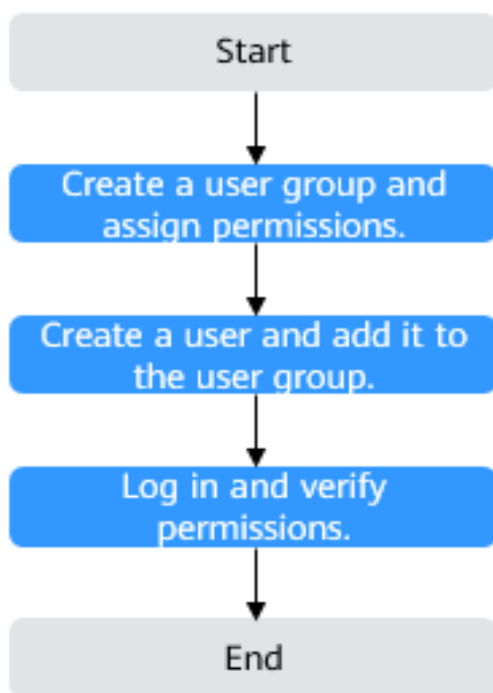
This section describes the procedure for granting permissions (see [Figure 1-2](#)).

Prerequisites

Learn about the permissions supported by MRS by referring to [Permission Management](#). For the permissions of other services, see [Permission Description](#).

Process Flow

Figure 1-2 Process for granting MRS permissions



1. [Creating a User Group and Assigning Permissions](#)

Create a user group on the IAM console, and assign MRS permissions to the group.

2. [Create a user and add it to a user group.](#)

Create a user on the IAM console and add the user to the group created in [1](#).

3. [Log in](#) and verify permissions.

Log in to the console by using the user created, and verify that the user has the granted permissions.

- Choose **Service List > Analytics > MapReduce Service**. Click **Buy Cluster** on the MRS console. If you fail to buy an MRS cluster (assume that you only have the **MRS ReadOnlyAccess** permission), the **MRS ReadOnlyAccess** policy has taken effect.

- Choose any other service in **Service List**. If a message appears indicating that you have insufficient permissions to access the service, the **MRS ReadOnlyAccess** policy has already taken effect.

IAM permissions of MRS

By default, new IAM users lack permissions assigned. Add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services.

MRS is a project-level service deployed for specific regions. To assign permissions to a user group, specify **Scope** as **Region-specific projects** and select projects in the corresponding region for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing MRS, the users need to switch to the authorized region.

You can grant permissions by using roles and policies.

- **Roles:** A coarse-grained IAM authorization strategy to assign permissions based on user responsibilities. Only a limited number of service-level roles are available. Some roles depend on other roles to take effect. When you assign such roles to users, remember to assign the roles they depend on. Roles are not an ideal choice for fine-grained authorization and secure access control.
- **Policies:** A fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This type of authorization is more flexible and is ideal for secure access control. For example, you can grant MRS users only the permissions for performing specified operations on MRS clusters, such as creating a cluster and querying a cluster list rather than deleting a cluster. Most policies define permissions based on APIs. For the API actions supported by MRS, see [Permissions Policies and Supported Actions](#).

Table 1-1 lists all the default system policies supported by MRS.

Table 1-1 MRS system policies

Policy	Description	Type
MRS FullAccess	Administrator permissions for MRS. Users with these permissions can perform all operations on MRS resources.	Fine-grained policy
MRS CommonOperations	Common user permissions for MRS. Users with these permissions can use MRS but cannot add or delete resources.	Fine-grained policy
MRS ReadOnlyAccess	Read-only permission for MRS. Users with these permissions can only view MRS resources.	Fine-grained policy

Policy	Description	Type
MRS Administrator	Operation permissions: <ul style="list-style-type: none"> • All operations on MRS • Permissions of the Tenant Guest and Server Administrator policies, which must also be granted to the users 	RBAC policy

Table 1-2 lists the common operations supported by each system-defined policy or role of MRS. Select the permissions as needed.

Table 1-2 Common operations supported by each system-defined policy

Operation	MRS FullAccess	MRS CommonOperations	MRS ReadOnlyAccess	MRS Administrator
Creating a cluster	√	x	x	√
Resizing a cluster	√	x	x	√
Upgrading node specifications	√	x	x	√
Deleting a cluster	√	x	x	√
Querying cluster details	√	√	√	√
Listing all clusters	√	√	√	√
Configuring an auto scaling rule	√	x	x	√
Querying a host list	√	√	√	√
Querying operation logs	√	√	√	√

Operation	MRS FullAccess	MRS CommonOperations	MRS ReadOnlyAccess	MRS Administrator
Creating and executing a job	√	√	x	√
Stopping a job	√	√	x	√
Deleting a single job	√	√	x	√
Deleting jobs in batches	√	√	x	√
Querying job details	√	√	√	√
Querying a job list	√	√	√	√
Creating a folder	√	√	x	√
Deleting a file	√	√	x	√
Querying a file list	√	√	√	√
Operating cluster tags in batches	√	√	x	√
Creating a single cluster tag	√	√	x	√
Deleting a single cluster tag	√	√	x	√
Querying a resource list by tag	√	√	√	√
Querying cluster tags	√	√	√	√
Accessing Manager	√	√	x	√

Operation	MRS FullAccess	MRS CommonOperations	MRS ReadOnlyAccess	MRS Administrator
Querying a patch list	√	√	√	√
Installing a patch	√	√	x	√
Uninstalling a patch	√	√	x	√
Authorizing O&M channels	√	√	x	√
Sharing O&M channel logs	√	√	x	√
Querying an alarm list	√	√	√	√
Subscribing to alarm notifications	√	√	x	√
Submitting a SQL statement	√	√	x	√
Querying SQL query results	√	√	x	√
Canceling a SQL execution task	√	√	x	√

1.3 Creating a Custom Policy for MRS

Custom policies can be created to supplement the system-defined policies of MRS. For the actions that can be added to custom policies, see [Permissions Policies and Supported Actions](#).

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.

- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see [Creating a Custom Policy](#).

 **NOTE**

Custom policy modifications do not take effect immediately. You need to wait about 15 minutes.

The following section contains examples of common MRS custom policies.

Example Custom Policies

- Example 1: Allowing users to create MRS clusters only

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "mrs:cluster:create",
        "ecs:*:*",
        "bms:*:*",
        "evs:*:*",
        "vpc:*:*",
        "smn:*:*"
      ]
    }
  ]
}
```

- Example 2: Allowing users to modify MRS clusters.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "mrs:cluster:resize"
      ]
    }
  ]
}
```

- Example 3: Allowing users to create a cluster, create and execute a job, and delete a single job, but denying cluster deletion

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "mrs:cluster:create",
        "mrs:job:submit",
        "mrs:job:delete"
      ]
    },
    {
      "Effect": "Deny",
      "Action": [
        "mrs:cluster:delete"
      ]
    }
  ]
}
```

- Example 4: Granting users the minimum permission to create an MRS cluster with ECS specifications

 NOTE

- If you need a key pair when creating a cluster, add the following permissions: **ecs:serverKeypairs:get** and **ecs:serverKeypairs:list**.
- Add the **kms:cmk:list** permission when encrypting data disks during cluster creation.
- Add the **mrs:alarm:subscribe** permission to enable the alarm function during cluster creation.
- Add the **rds:instance:list** permission to use external data sources during cluster creation.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "mrs:cluster:create"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "ecs:cloudServers:updateMetadata",
        "ecs:cloudServerFlavors:get",
        "ecs:cloudServerQuotas:get",
        "ecs:servers:list",
        "ecs:servers:get",
        "ecs:cloudServers:delete",
        "ecs:cloudServers:list",
        "ecs:serverInterfaces:get",
        "ecs:serverGroups:manage",
        "ecs:servers:setMetadata",
        "ecs:cloudServers:get",
        "ecs:cloudServers:create"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "vpc:securityGroups:create",
        "vpc:securityGroupRules:delete",
        "vpc:vpcs:create",
        "vpc:ports:create",
        "vpc:securityGroups:get",
        "vpc:subnets:create",
        "vpc:privateIps:delete",
        "vpc:quotas:list",
        "vpc:networks:get",
        "vpc:publicIps:list",
        "vpc:securityGroups:delete",
        "vpc:securityGroupRules:create",
        "vpc:privateIps:create",
        "vpc:ports:get",
        "vpc:ports:delete",
        "vpc:publicIps:update",
        "vpc:subnets:get",
        "vpc:publicIps:get",
        "vpc:ports:update",
        "vpc:vpcs:list"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "evs:quotas:get",
        "evs:types:get"
      ]
    }
  ]
}
```



```
    },  
    {  
      "Effect": "Allow",  
      "Action": [  
        "bms:serverFlavors:get"  
      ]  
    }  
  ]  
}
```

- Example 5: Granting users the minimum permission to create an MRS cluster with BMS specifications

NOTE

- If you need a key pair when creating a cluster, add the following permissions: **ecs:serverKeypairs:get** and **ecs:serverKeypairs:list**.
- Add the **kms:cmk:list** permission when encrypting data disks during cluster creation.
- Add the **mrs:alarm:subscribe** permission to enable the alarm function during cluster creation.
- Add the **rds:instance:list** permission to use external data sources during cluster creation.

```
{  
  "Version": "1.1",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "mrs:cluster:create"  
      ]  
    },  
    {  
      "Effect": "Allow",  
      "Action": [  
        "ecs:servers:list",  
        "ecs:servers:get",  
        "ecs:cloudServers:delete",  
        "ecs:serverInterfaces:get",  
        "ecs:serverGroups:manage",  
        "ecs:servers:setMetadata",  
        "ecs:cloudServers:create",  
        "ecs:cloudServerFlavors:get",  
        "ecs:cloudServerQuotas:get"  
      ]  
    },  
    {  
      "Effect": "Allow",  
      "Action": [  
        "vpc:securityGroups:create",  
        "vpc:securityGroupRules:delete",  
        "vpc:vpcs:create",  
        "vpc:ports:create",  
        "vpc:securityGroups:get",  
        "vpc:subnets:create",  
        "vpc:privateIps:delete",  
        "vpc:quotas:list",  
        "vpc:networks:get",  
        "vpc:publicIps:list",  
        "vpc:securityGroups:delete",  
        "vpc:securityGroupRules:create",  
        "vpc:privateIps:create",  
        "vpc:ports:get",  
        "vpc:ports:delete",  
        "vpc:publicIps:update",  
        "vpc:subnets:get",  
        "vpc:publicIps:get",  
      ]  
    }  
  ]  
}
```

```
        "vpc:ports:update",
        "vpc:vpcs:list"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "evs:quotas:get",
      "evs:types:get"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "bms:servers:get",
      "bms:servers:list",
      "bms:serverQuotas:get",
      "bms:servers:updateMetadata",
      "bms:serverFlavors:get"
    ]
  }
]
}
```

- Example 6: Allowing users to create a hybrid ECS and BMS cluster with the minimum permission

NOTE

- If you need a key pair when creating a cluster, add the following permissions: **ecs:serverKeyPairs:get** and **ecs:serverKeyPairs:list**.
- Add the **kms:cmk:list** permission when encrypting data disks during cluster creation.
- Add the **mrs:alarm:subscribe** permission to enable the alarm function during cluster creation.
- Add the **rds:instance:list** permission to use external data sources during cluster creation.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "mrs:cluster:create"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "ecs:cloudServers:updateMetadata",
        "ecs:cloudServerFlavors:get",
        "ecs:cloudServerQuotas:get",
        "ecs:servers:list",
        "ecs:servers:get",
        "ecs:cloudServers:delete",
        "ecs:cloudServers:list",
        "ecs:serverInterfaces:get",
        "ecs:serverGroups:manage",
        "ecs:servers:setMetadata",
        "ecs:cloudServers:get",
        "ecs:cloudServers:create"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "vpc:securityGroups:create",

```

```
        "vpc:securityGroupRules:delete",
        "vpc:vpcs:create",
        "vpc:ports:create",
        "vpc:securityGroups:get",
        "vpc:subnets:create",
        "vpc:privateIps:delete",
        "vpc:quotas:list",
        "vpc:networks:get",
        "vpc:publicIps:list",
        "vpc:securityGroups:delete",
        "vpc:securityGroupRules:create",
        "vpc:privateIps:create",
        "vpc:ports:get",
        "vpc:ports:delete",
        "vpc:publicIps:update",
        "vpc:subnets:get",
        "vpc:publicIps:get",
        "vpc:ports:update",
        "vpc:vpcs:list"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "evs:quotas:get",
        "evs:types:get"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "bms:servers:get",
        "bms:servers:list",
        "bms:serverQuotas:get",
        "bms:servers:updateMetadata",
        "bms:serverFlavors:get"
    ]
}
]
```

2 MRS Cluster Planning

2.1 Service Selection

2.1.1 MRS Cluster Types

MRS consists of multiple big data components, and you can select the cluster type that best fits your service requirements, data types, reliability expectations, and resource budget.

You can **quickly buy** a cluster using the preset cluster template or select the component list and advanced settings to **manually buy** a cluster.

Table 2-1 MRS cluster types

Type	Scenario	Core component
Hadoop analysis cluster	Hadoop cluster uses components in the open source Hadoop ecosystem to analyze and query vast amounts of data. For example, use YARN to manage cluster resources, Hive and Spark to provide offline storage and computing of large-scale distributed data, Spark Streaming and Flink to offer streaming data computing, and Tez to provide a distributed computing framework of directed acyclic graphs (DAGs).	Hadoop, Hive, Spark, Tez, Flink, ZooKeeper, and Ranger

Type	Scenario	Core component
HBase query cluster	An HBase cluster uses Hadoop and HBase components to provide a column-oriented distributed cloud storage system featuring enhanced reliability, great performance, and elastic scalability. It applies to the storage and distributed computing of massive amounts of data. You can use HBase to build a storage system capable of storing TB- or even PB-level data. With HBase, you can filter and analyze data with ease and get responses in milliseconds, rapidly mining data value.	Hadoop, HBase, ZooKeeper, and Ranger
Kafka streaming cluster	Kafka cluster uses Kafka and Storm to provide an open source message system with high throughput and scalability. It is widely used in scenarios such as log collection and monitoring data aggregation to implement efficient streaming data collection and real-time data processing and storage.	Kafka and Storm
ClickHouse cluster	ClickHouse is a columnar database management system for online analysis. It features ultimate compression ratio and fast query performance. It is widely used in Internet advertisement, app and web traffic analysis, telecom, finance, and IoT fields.	ClickHouse and ZooKeeper
Real-time analysis cluster	Real-time analysis clusters use Hadoop, Kafka, Flink, and ClickHouse components to provide a system for collection, real-time analysis, and query of data at scale.	Hadoop, Kafka, Flink, ClickHouse, ZooKeeper, and Ranger

2.1.2 MRS Cluster Node Types

An MRS cluster consists of multiple ECSs. The system manages nodes in node groups based on specifications. Nodes in the same node group use same ECS specifications.

Nodes in a cluster can be classified into Master nodes, Core nodes, and Task nodes based on the roles of components deployed on the nodes.

Table 2-2 Cluster node types

Node Type	Functions
Master node	<p>MRS cluster management node. OMS server is deployed on the node to manage and monitor the cluster.</p> <p>After creating the MRS cluster, the node containing master1 in its name functions as the Master1 node, and the one with master2 in its name functions as the Master2 node.</p> <p>You can log in to a Master node either using VNC on the ECS management console or using SSH. After logging in to the Master node, you can access other nodes without entering passwords.</p> <p>The system automatically deploys the Master nodes in active/standby mode and supports the high availability (HA) feature for MRS cluster management. If the active management node fails, the standby management node switches to the active state and takes over services.</p> <p>To determine whether the Master1 node is the active management node, see Checking MRS Active/Standby Management Nodes.</p>
Core node	<p>Work node of an MRS cluster. It processes and analyzes data and stores process data.</p> <p>In the Nodes tab of the cluster details page, the nodes in the node group whose Node Type includes Core are Core nodes.</p>
Task node	<p>Compute node. When the compute resources of a cluster are insufficient, you can configure elastic scaling policies to increase nodes automatically.</p> <p>In the Nodes tab of the cluster details page, the nodes in the node group whose Node Type is Task are Task nodes.</p> <p>If only the NodeManager (Yarn) or Supervisor (Storm) role is deployed in a node group in addition to basic mandatory roles, this node group is a Task node group.</p>

MRS cluster nodes support remote login. The following remote login methods are available:

- GUI login: Use the remote login function provided by the ECS console to log in to the Linux CLI of the cluster node.
- SSH login: Use a remote login tool like PuTTY to access Linux ECSs. You need to bind an EIP to the ECS first.

Apply for and bind an EIP to a cluster node. For details, see [Assigning an EIP](#).

You can log in to the ECS using a key or password. For details, see [Logging In to an MRS Cluster Node](#).

2.1.3 MRS Cluster Node Specifications

MRS Node Specifications

MRS supports host specifications determined by CPU, memory, and disk space.

Tenants share physical resources of ECSs, but can exclusively use resources of BMSs. BMSs can better meet your requirements for deploying key applications and services that require high performance (such as big data clusters and enterprise middleware systems) and a secure and reliable running environment. If BMS specifications are used, Master node specifications cannot be scaled up.

MRS supports BMS specifications only when the billing mode of a cluster is **Yearly/Monthly**.

MRS supports the following hybrid deployment of **ECSs** and **BMSs**:

- Master, Core, and Task nodes are deployed on ECSs.
- Master and Core nodes are deployed on BMSs, and Task nodes are deployed on ECSs.
- Master nodes are deployed on either ECSs or BMSs, Core nodes are deployed on either ECSs or BMSs, and Task nodes are deployed on ECSs.

Tenants share physical resources of ECSs, but can exclusively use resources of BMSs. BMSs can better meet your requirements for deploying key applications and services that require high performance (such as big data clusters and enterprise middleware systems) and a secure and reliable running environment.

If BMS specifications are used, Master node specifications cannot be scaled up.

NOTE

- More advanced instance specifications provide better data processing. However, they require higher cluster cost.
- Instance specifications may vary in different AZs. If no instance specifications in the current AZ can meet your requirements, switch to another AZ.
- If you select HDDs for Core nodes, there is no billing information for data disks. The fees are charged with ECSs.
- If you select non-HDD disks for Core nodes, the disk types of Master and Core nodes are determined by **Data Disk**.
- If **Sold out** appears next to an instance specification of a node, the node of this specification cannot be bought. You can only buy nodes of other specifications.
- The Master node specification (4 vCPUs 8 GB) is not within the SLA after-sales scope. It is applicable only to the test environment and is not recommended for the production environment.
- For MRS 3.x or later, the memory of the master node must be greater than 64 GB.

Disk Roles

Table 2-3 Disk types of MRS cluster nodes

Disk Role	Description
System disk	<p>Storage type and space of the system disk on a node.</p> <p>Storage type can be any of the following:</p> <ul style="list-style-type: none">• SAS: high I/O• SSD: ultra-high I/O• GPSSD: general-purpose SSD
Data disk	<p>Data disk storage space of a node. For more data storage, you can add disks when creating a cluster. A maximum of 10 disks can be added to each Core or Task node.</p> <ul style="list-style-type: none">• Data storage and computing are separated. Data is stored in OBS, which features low cost and unlimited storage capacity. The clusters can be deleted at any time in OBS. The computing performance is determined by OBS access performance and is lower than that of HDFS. This configuration is recommended if data computing is infrequent.• Data storage and computing are not separated. Data is stored in HDFS, which features high cost, high computing performance, and limited storage capacity. Before deleting clusters, you must export and store the data. This configuration is recommended if data computing is frequent. <p>The storage type can be any of the following:</p> <ul style="list-style-type: none">• SAS: high I/O• SSD: ultra-high I/O• GPSSD: general-purpose SSD <p>NOTE</p> <p>Adding nodes to an MRS cluster requires increasing the disk capacity of the management node (Master node). To ensure stable cluster running, set the disk capacity of the Master node to over 600 GB if the number of nodes is 300 and increase it to over 1 TB if the number of nodes reaches 500.</p>

2.2 MRS Cluster Deployment

2.2.1 Overview

The analysis cluster, streaming cluster, and hybrid cluster provided by MRS use fixed templates to deploy cluster processes. Therefore, you cannot customize service processes on management nodes and control nodes.

If you want to customize the cluster deployment, set **Cluster Type** to **Custom** when creating a cluster. In this way, you can customize the deployment mode of process instances on the management nodes and control nodes in the cluster.

 **NOTE**

Only MRS 3.x and later versions support the creation of clusters in a custom topology.

A custom cluster provides the following functions:

- Separated deployment of the management and control roles: The management role and control role are deployed on different Master nodes.
- Co-deployment of the management and control roles: The management and control roles are co-deployed on the Master node.
- Components are deployed separately to avoid resource contention.

MRS Cluster Deployment Types

Table 2-4 MRS cluster deployment types

Common Node	Description	Node Range
Compact	The management role and control role are deployed on the Master node, and data instances are deployed in the same node group. This deployment mode applies to scenarios where the number of control nodes is less than 100, reducing costs.	<ul style="list-style-type: none">• The number of Master nodes is greater than or equal to 3 and less than or equal to 11.• The total number of node groups is less than or equal to 10, and the total number of nodes in non-Master node groups is less than or equal to 10,000.
Co-deployment of the management role, control role, and data instances	This template is not recommended in the production environment or commercial environment. If management, control, and data nodes are deployed together, cluster performance and reliability may deteriorate. If there are enough nodes, deploy data nodes separately.	<ul style="list-style-type: none">• There must be at least 3 master nodes and 100 at most.• The total node groups are no more than 10.
OMS-separate	The management role and control role are deployed on different Master nodes, and data instances are deployed in the same node group. This deployment mode is applicable to a cluster with 100 to 500 nodes and delivers better performance in high-concurrency load scenarios.	<ul style="list-style-type: none">• The number of Master nodes is greater than or equal to 5 and less than or equal to 11.• The total number of node groups is less than or equal to 10, and the total number of nodes in non-Master node groups is less than or equal to 10,000.

Common Node	Description	Node Range
Full-size	The management role and control role are deployed on different Master nodes, and data instances are deployed in different node groups. This deployment mode is applicable to a cluster with more than 500 nodes. Components can be deployed separately, which can be used for a larger cluster scale.	<ul style="list-style-type: none"> The number of Master nodes is greater than or equal to 9 and less than or equal to 11. The total number of node groups is less than or equal to 10, and the total number of nodes in non-Master node groups is less than or equal to 10,000.

Role Deployment Rules for MRS Clusters

The MRS system consists of multiple services in logical architecture. Each service contains one or more roles. Each role can deploy one or more instances.

- **Service:** A service is a kind of service capabilities provided by a component in a cluster. Each component in the cluster provides a service and is named by the service.
- **Role:** Roles are components of a service. Each service consists of one or more roles. Services are deployed on nodes (servers) through roles to ensure normal service running.
- **Instance:** An instance is formed when a service role is installed on a node. A service has one or more role instances.

For details about the deployment principles of each service, see [Overview](#).

- **Service A depending on service B:** If service A is deployed in a cluster, service B must be deployed in advance. Service B provides basic capabilities for service A.
- **Service A associated with service B:** Service A exchanges data with service B during service running but does not depend on service B during deployment.
- **Role A and role B deployed on the same server:** If role A is deployed in a cluster, role B must also be deployed, and role A and role B must be deployed on the same node.

Table 2-5 MRS role deployment rules

Service	Dependency	Role	Role Deployment Suggestions
OMSServer	-	OMSServer	This role can be deployed it on the Master node and cannot be modified.

Service	Dependency	Role	Role Deployment Suggestions
ClickHouse	Depends on ZooKeeper.	CHS (ClickHouseServer)	This role can be deployed on all nodes. You can deploy 2 to 256 role instances, but the number must be even.
		CLB (ClickHouseBalancer)	This role can be deployed on all nodes. Number of role instances to be deployed: 2 to 256
Flink	<ul style="list-style-type: none"> • Depends on ZooKeeper. • Depends on KrbServer. • Depends on DBService. • Depends on Hadoop. 	FR(FlinkResource)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000
		FS(FlinkServer)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 2
Flume	-	MS(MonitorServer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 2
		F(Flume)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000
Hadoop	Depends on ZooKeeper.	NN(NameNode)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2
		HFS (HttpFS)	This role can be deployed on the Master node only. Number of role instances to be deployed: 0 to 10
		JN(JournalNode)	This role can be deployed on the Master node only. Number of role instances to be deployed: 3 to 60, with the step size of 2

Service	Dependency	Role	Role Deployment Suggestions
		DN(DataNode)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000
		RM(Resource Manager)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2
		NM(NodeManager)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000
		JHS(JobHistoryServer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 2
		TLS(Timeline Server)	This role can be deployed on the Master node only. Number of role instances to be deployed: 0 to 1
HBase	Depends on Hadoop.	HM(HMaster)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2
		TS(ThriftServer)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 10,000
		RT(RESTServer)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 10,000
		RS(RegionServer)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000

Service	Dependency	Role	Role Deployment Suggestions
		TS1(Thrift1Server)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 10,000 If the Hue service is installed in a cluster and HBase needs to be used on the Hue web UI, install this instance for HBase.
HetuEngine	<ul style="list-style-type: none"> • Depends on Hadoop. • Depends on DBService. • Depends on Hive. • Depends on ZooKeeper. • Depends on KrbServer. • Depends on Yarn. • Depends on HDFS. 	HSB(HSBroker)	This role can be deployed on all nodes. Number of role instances to be deployed: 2 to 50
		HSC(HSConsole)	This role can be deployed on all nodes. Number of role instances to be deployed: 2
		HSF(HSFabric)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 50
		QAS (available for MRS 3.2.0-LTS.1 and later versions only)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 2
Hive	<ul style="list-style-type: none"> • Depends on Hadoop. • Depends on DBService. 	MS(MetaStore)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2 to 10
		WH (WebHCat)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 10
		HS(HiveServer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2 to 80
Hue	Depends on DBService.	H(Hue)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2

Service	Dependency	Role	Role Deployment Suggestions
Impala	<ul style="list-style-type: none"> • Depends on Hadoop. • Depends on Hive. • Depends on DBService. • Depends on ZooKeeper. 	StateStore	This role can be deployed on the Master node only. Number of role instances to be deployed: 1
		Catalog	This role can be deployed on the Master node only. Number of role instances to be deployed: 1
		Impalad	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000
IoTDB	Depends on KrbServer.	ConfigNode (CN)	This role can be deployed on Master nodes only. Number of role instances to be deployed: 3 to 9, with the step size of 2
		IoTDBServer (IoTDBS)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 256
Kafka	Depends on ZooKeeper.	B(Broker)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000
Kudu	-	KuduMaster	This role can be deployed on the Master node only. Number of role instances to be deployed: 3 or 5
		KuduTserver	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000
Loader	<ul style="list-style-type: none"> • Depends on Hadoop. • Depends on DBService. 	LS(LoaderServer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2

Service	Dependency	Role	Role Deployment Suggestions
Oozie	<ul style="list-style-type: none"> • Depends on Hadoop. • Depends on DBService. • Depends on ZooKeeper. 	O(oozie)	<p>This role can be deployed on the Master node only.</p> <p>Number of role instances to be deployed: 2</p>
Presto	Depends on Hive.	PCD(Coordinator)	<p>This role can be deployed on the Master node only.</p> <p>Number of role instances to be deployed: 2</p>
		PWK(Worker)	<p>This role can be deployed on all nodes.</p> <p>Number of role instances to be deployed: 1 to 10,000</p>
Ranger	Depends on DBService.	RA(RangerAdmin)	<p>This role can be deployed on the Master node only.</p> <p>Number of role instances to be deployed: 1 to 2</p>
		USC(UserSync)	<p>This role can be deployed on the Master node only.</p> <p>Number of role instances to be deployed: 1</p>
		TSC(TagSync)	<p>This role can be deployed on all nodes.</p> <p>Number of role instances to be deployed: 0 to 1</p>
Spark	<ul style="list-style-type: none"> • Depends on Hadoop. • Depends on Hive. • Depends on ZooKeeper. 	JS(JDBCServer)	<p>This role can be deployed on the Master node only.</p> <p>Number of role instances to be deployed: 1 to 2</p>
		JH(JobHistory)	<p>This role can be deployed on the Master node only.</p> <p>Number of role instances to be deployed: 1 to 2</p>
		SR(SparkResource)	<p>This role can be deployed on the Master node only.</p> <p>Number of role instances to be deployed: 1 to 2</p>

Service	Dependency	Role	Role Deployment Suggestions
Spark2x	<ul style="list-style-type: none"> • Depends on Hadoop. • Depends on Hive. • Depends on ZooKeeper. 	JS2X(JDBCServer2x)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2 to 10
		JH2X(JobHistory2x)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2
		SR2X(SparkResource2x)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2 to 50
		IS2X(IndexServer2x)	(Optional) This role can be deployed on the Master node only. Number of role instances to be deployed: 0 to 2, with the step size of 2
Sqoop	Depends on Hadoop.	SC(SqoopClient)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000
Tez	<ul style="list-style-type: none"> • Depends on Hadoop. • Depends on DBService. • Depends on ZooKeeper. 	TUI(TezUI)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 2
ZooKeeper	-	QP(quorumpeer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 3 to 9, with the step size of 2

2.2.2 Kerberos Authentication for MRS Clusters

The Hadoop community version has two authentication modes: Kerberos authentication for security and Simple authentication for regular use. During cluster setup, you can choose to enable or disable Kerberos authentication. Once the cluster is created, it cannot be modified.

Security Mode (Kerberos Authentication Enabled)

The MRS clusters in security mode use the Kerberos authentication protocol for security authentication. The Kerberos protocol supports mutual authentication between clients and servers. This eliminates the risks incurred by sending user credentials over the network for simulated authentication. In clusters, KrbServer provides the Kerberos authentication support.

Kerberos user object

In the Kerberos protocol, each user object is a principal. A complete principal consists of username and domain name. In O&M or application development scenarios, the user identity must be verified before a client connects to a server. Users for O&M and service operations are classified into human-machine and machine-machine users. The password of human-machine users is manually configured, while the password of machine-machine users is generated by the system randomly.

Kerberos authentication

Kerberos authentication supports two authentication modes: password authentication and keytab authentication. The default authentication validity period is 24 hours.

- Password authentication: User identity is verified by entering the correct password. This mode mainly used in O&M scenarios where human-machine users are used. The client command is **kinit Username**.
- Keytab authentication: Keytab files contain users' principal and encrypted credential information. When keytab files are used for authentication, the system automatically uses encrypted credential information to perform authentication and the user password does not need to be entered. This mode is mainly used in component application development scenarios where machine-machine users are used. Keytab authentication can also be configured using the **kinit** command.

Common mode (Kerberos Authentication Disabled)

In normal clusters, MRS components use a native open source authentication mechanism, which is typically Simple authentication. If Simple authentication is used, authentication is automatically performed by a client user (for example, user **root**) by default when a client connects to a server. The authentication is imperceptible to the administrator or service user. In addition, when being executed, the client may even pretend to be any user (including **superuser**) by injecting **UserGroupInformation**. Cluster resource management and data control APIs are not authenticated on the server and are easily exploited and attacked by hackers.

It is best to use this for single-user scenarios in normal clusters, but make sure to tightly control network access permissions to keep the cluster secure.

- Deploy service applications on ECSs in the same VPC and subnet to avoid accessing the MRS cluster through the external network.
- Configure security group rules to strictly control the access scope. Do not configure access rules that allow **Any** or **0.0.0.0** for the inbound direction of MRS cluster ports.

2.2.3 ECS Specifications Supported by MRS Clusters

MRS uses ECSs of the following types in different application scenarios.

- General computing-plus: C3, C3ne, C6, and C6s
- Memory-optimized: M3, and M6
- Ultra-high I/O: I3 and IR3
- Kunpeng general computing-plus: KC1

MRS uses the following BMS:

- Kunpeng V1 BMS

ECS Flavor Naming Rules

AB.C.D

Example: m2.8xlarge.8

In the preceding flavor:

- **A** specifies the ECS type. For example, **s** indicates a general-purpose ECS, **c** a computing ECS, and **m** a memory-optimized ECS.
- **B** specifies the type ID. For example, the **1** in **s1** indicates a general-purpose first-generation ECS, and the **2** in **s2** indicates a general-purpose second-generation ECS.
- **C** specifies a flavor size and can be any of the following options: medium, large, and xlarge.
- **D** specifies the ratio of memory to vCPUs expressed in a digit. For example, value **4** indicates that the ratio of memory to vCPUs is 4.

ECS Specifications

Table 2-6 General computing-plus (C) ECS specifications

Type	vCPU	Memory (GB)	Flavor	Virtualization Type
C3	32	64	c3.8xlarge.2	KVM
C3	16	64	c3.4xlarge.4	KVM
C3	32	128	c3.8xlarge.4	KVM
C3	60	256	c3ne.15xlarge.4	KVM
C3ne	32	64	c3ne.8xlarge.2	KVM
C3ne	16	64	c3ne.4xlarge.4	KVM
C3ne	32	128	c3ne.8xlarge.4	KVM
C3ne	60	256	c3ne.15xlarge.4	KVM

Type	vCPU	Memory (GB)	Flavor	Virtualization Type
C6	32	64	c6.8xlarge.2	KVM
C6	64	128	c6.16xlarge.2	KVM
C6	16	64	c6.4xlarge.4	KVM
C6	32	128	c6.8xlarge.4	KVM
C6	64	256	c6.16xlarge.4	KVM
C6s	32	64	c6s.8xlarge.2	KVM
C6s	64	128	c6s.16xlarge.2	KVM

Table 2-7 Memory-optimized ECS specifications

Type	vCPU	Memory (GB)	Flavor	Virtualization Type
M3	8	64	m3.2xlarge.8	KVM
M3	16	128	m3.4xlarge.8	KVM
M3	32	256	m3.8xlarge.8	KVM
M3	60	512	m3.15xlarge.8	KVM
M6	8	64	m6.2xlarge.8	KVM
M6	16	128	m6.4xlarge.8	KVM
M6	32	256	m6.8xlarge.8	KVM
M6	64	512	m6.16xlarge.8	KVM

Table 2-8 Ultra-high I/O ECS specifications

Type	vCPU	Memory (GB)	Flavor	Virtualization Type
I3	8	64	i3.2xlarge.8	KVM
I3	16	128	i3.4xlarge.8	KVM
I3	32	256	i3.8xlarge.8	KVM
I3	64	512	i3.16xlarge.8	KVM
IR3	16	64	ir3.4xlarge.4	KVM

Type	vCPU	Memory (GB)	Flavor	Virtualization Type
IR3	32	128	ir3.8xlarge.4	KVM

BMS Specifications

Table 2-9 Kunpeng V1 BMS specifications

Flavor/ID	vCPU	Memory (GB)	Network
physical.ks1ne.4xlarge	128	512	Distributed
physical.ks1ne.8xlarge	128	1024	

3 Buying MRS Clusters

3.1 Quickly Buying an MRS Cluster

MRS consists of multiple big data components, and you can select the cluster type that best fits your service requirements, data types, reliability expectations, and resource budget.

This section uses an HBase query cluster as an example to describe how to quickly purchase an MRS cluster. An HBase cluster uses Hadoop and HBase components to provide a column-oriented distributed cloud storage system featuring enhanced reliability, great performance, and elastic scalability. It applies to the storage and distributed computing of massive amounts of data. You can use HBase to build a storage system capable of storing TB- or even PB-level data. With HBase, you can filter and analyze data with ease and get responses in milliseconds, rapidly mining data value.

Procedure

- Step 1** Go to the [Buy Cluster](#) page.
- Step 2** On the displayed page, click the **Quick Config** tab.
- Step 3** Configure basic cluster information by referring to the following table.

Table 3-1 MRS cluster parameters

Parameter	Description	Example Value
Billing Mode	Billing mode of a cluster. MRS provides two billing modes: yearly/monthly and pay-per-use. Pay-per-use: If you select this mode, a prepaid balance will be frozen. For details, see Billing Description .	Pay-per-use

Parameter	Description	Example Value
Region	Region where the resource to be created is located. Resources located in different regions cannot communicate with each other using an internal network. To improve access speed and reduce network latency, choose a region that is close to your own.	-
Cluster Name	MRS cluster name. You can use the default name. However, you are advised to include a project name abbreviation or date for consolidated memory and easy distinguishing. After a cluster is created, you can change the cluster name in the cluster list.	mrs-test
Cluster Type	Select a proper MRS cluster type based on service requirements. <ul style="list-style-type: none"> • Analysis cluster: It is used for offline data analysis and comprises data analysis tools like Hadoop, Spark, HBase, Hive, Flink, Oozie, and Tez. • Streaming cluster: It processes streaming data to quickly analyze real-time data sources, and it mainly includes streaming data processing tools such as Kafka and Flume. • Hybrid cluster: It is utilized for both offline data analysis and stream processing. • Custom: You can select from a variety of components that are supported by the corresponding version of the MRS cluster. 	Custom

Parameter	Description	Example Value
Version Type	<p>MRS provides two types of clusters: LTS and Normal. Different versions provide different components. You can select a version as required.</p> <ul style="list-style-type: none"> • LTS: employs MRS's own components to provide highly reliable clusters with strong DR capabilities, making long-term support and evolution possible. • Normal: integrates MRS's mature and stable features and functions with open-source capabilities, offering high performance and stability. 	LTS
Cluster Version	Version of the MRS cluster. Different versions may contain different open source component versions and functions. You are advised to select the latest version.	MRS 3.2.0-LTS.1
Component	Cluster templates containing preset opensource components you will need for your business.	HBase Query Cluster
AZ	Resources are assigned to the AZ of the current region during creation. An AZ is a physical region that operates on independent power supplies and networks for resource usage.	AZ1
VPC	VPC to which the MRS cluster node belongs. If no VPC is available, click View VPC to access the network console and create a VPC.	-
Subnet	Subnet information in the VPC. If no subnet is available, click View Subnet to access the network console and create a subnet.	-
Cluster Node	<p>Specifications and quantity of nodes in an MRS cluster.</p> <p>For MRS 3.x or later, the memory of the master node must be greater than 64 GB.</p>	Select the number of cluster node specifications as required.

Parameter	Description	Example Value
Kerberos Authentication	Whether to enable Kerberos authentication for each component in the MRS cluster. If Kerberos authentication is enabled, users can access component resources only after being authenticated. This function cannot be changed after you buy a cluster.	Switch this function on.
Username	Default user for logging in to Manager and nodes in the MRS cluster. User admin is used to log in to Manager, while user root is used to log in to the OS of nodes in the cluster.	-
Password/Confirm Password	Set the passwords for the root and admin users. The passwords are user-defined and must be kept secure.	-
Enterprise Project	Enterprise project is a way to manage cloud resources. It allows you to manage resources and members within a project, and you can choose to use the system-defined enterprise project default or create your own.	default
Secure Communications	To allow the MRS console to access big data components in the user VPC, you need to activate the relevant security group rules. For details, see Configuring Secure Communication Authorization for an MRS Cluster .	Select this function.

Step 4 Click **Buy Now**.

If Kerberos authentication is enabled, check whether this function is required. If it is, click **Continue**. If not, click **Back** to disable it and then proceed with the subsequent step. This function cannot be changed after you buy a cluster.

NOTE

For any doubt about the pricing, click **Pricing details** in the lower left corner.

Step 5 Click **Back to Cluster List** to view the cluster status. Click **Access Cluster** to view cluster details.

For details about cluster status during creation, see the description of the status parameters in [Table 6-9](#).

It takes some time to create a cluster. The initial status of the cluster is **Starting**. After the cluster has been created successfully, the cluster status becomes **Running**.

On the MRS management console, a maximum of 10 clusters can be concurrently created, and a maximum of 100 clusters can be managed.

----End

3.2 Manually Buying an MRS Cluster

This section describes how to create an MRS cluster on the MRS console.

MRS consists of multiple big data components, and you can select the cluster type that best fits your service requirements, data types, reliability expectations, and resource budget.

You can **quickly buy** a cluster using the preset cluster template or select the component list and advanced settings to manually buy a cluster.

Manually Buying an MRS Cluster

Step 1 Go to the **Buy Cluster** page.

Step 2 Click the **Custom Config** tab.

NOTE

When creating a cluster, pay attention to quota notification. If a resource quota is insufficient, increase the resource quota as prompted and create a cluster.

Step 3 In the basic configuration area, set basic information about the MRS cluster.

- **Billing Mode:** MRS provides two billing modes: yearly/monthly and pay-per-use.
- **Region:** Resources in different regions cannot communicate through intranet. To enhance access speed and minimize network latency, it is recommended to choose the nearest region when selecting the location of your resources.

Step 4 Configure the MRS cluster information as prompted.

When creating a cluster, you need to configure the cluster version, network, node specifications and quantity, and other advanced settings.

- **MRS Cluster Version Configuration**
- **MRS Cluster Network Configuration**
- **MRS Cluster Node Configuration**
- **Other MRS Cluster Configuration Parameters**

Step 5 After the cluster information is configured, click **Buy Now**.

If Kerberos authentication is enabled for a cluster, check whether Kerberos authentication is required. If yes, click **Continue**. If no, click **Back** to disable Kerberos authentication and then create a cluster. This function cannot be changed after you buy a cluster.

NOTE

- For any doubt about the pricing, click **Pricing details** in the lower left corner.
- If you select the pay-per-use billing mode, the order may fail to be placed because the account may be in risk. In this case, contact technical support.

Step 6 Click **Back to Cluster List** to view the cluster status.

Wait for the cluster creation to complete. The initial status of the cluster is **Starting**. After the cluster is created, the cluster status becomes **Running**.

On the MRS console, a maximum of 10 clusters can be created concurrently, and a maximum of 100 clusters can be managed.

For details about cluster status during creation, see the description of the status parameters in [Table 6-9](#).

----End

MRS Cluster Version Configuration

Table 3-2 MRS cluster parameters

Parameter	Description	Example Value
Cluster Type	<p>Select a proper MRS cluster type based on service requirements.</p> <ul style="list-style-type: none"> • Analysis cluster: It is used for offline data analysis and comprises data analysis tools like Hadoop, Spark, HBase, Hive, Flink, Oozie, and Tez. • Streaming cluster: It processes streaming data to quickly analyze real-time data sources, and it mainly includes streaming data processing tools such as Kafka and Flume. • Hybrid cluster: It is utilized for both offline data analysis and stream processing. • Custom: You can select from a variety of components that are supported by the corresponding version of the MRS cluster. 	Custom
Version Type	<p>MRS provides two types of clusters: LTS and Normal. Different versions provide different components. You can select a version as required.</p> <ul style="list-style-type: none"> • LTS: employs MRS's own components to provide highly reliable clusters with strong DR capabilities, making long-term support and evolution possible. • Normal: integrates MRS's mature and stable features and functions with open-source capabilities, offering high performance and stability. 	LTS
Cluster Version	<p>Version of the MRS cluster. Different versions may contain different open source component versions and functions. You are advised to select the latest version.</p>	MRS 3.2.0-LTS.1

Parameter	Description	Example Value
Component	<p>Select the components to be deployed in the MRS cluster.</p> <p>You can change components based on your needs. For some clusters, components cannot be added after creation.</p>	-
Metadata	<p>Whether to use external data sources to store Hive and Ranger metadata of the cluster.</p> <ul style="list-style-type: none"> • Local: Metadata is stored in the local cluster. • External data connection: Metadata of external data sources is used. If the cluster is abnormal or deleted, metadata is not affected. This mode applies to scenarios where storage and compute are decoupled. When creating an MRS cluster, you can connect to LakeFormation instances to store metadata of components such as Hive and Spark. For details, see Configuring a LakeFormation Data Connection. <p>After a cluster is created, you can manually store the component metadata to external data sources. For details, see Managing MRS Cluster Metadata.</p>	Local
Component Port	<p>Default communication port policy of each component in the MRS cluster. The LTS cluster supports the configuration.</p> <ul style="list-style-type: none"> • Open source: Use the port provided by the open source component. • Custom: Customize a port for the component. <p>For details about the differences between default open source port and default custom port, see Common Ports for MRS Cluster Services.</p>	Open source

MRS Cluster Network Configuration

Table 3-3 Network configuration parameters

Parameter	Description	Example Value
AZ	<p>Resources are assigned to the AZ of the current region during creation. An AZ is a physical region that operates on independent power supplies and networks for resource usage.</p>	AZ1

Parameter	Description	Example Value
VPC	VPC to which the MRS cluster node belongs. If no VPC is available, click View VPC to access the network console and create a VPC.	-
Subnet	<p>Subnet information in the VPC. If no subnet is available, click View Subnet to access the network console and create a subnet.</p> <p>A subnet provides dedicated network resources that are logically isolated from other networks for network security. For details about how to configure network ACL outbound rules, see How Do I Configure a Network ACL Outbound Rule?</p> <p>NOTE</p> <ul style="list-style-type: none">The number of IP addresses required by creating an MRS cluster depends on the number of cluster nodes and selected components, but not the cluster type.In MRS, IP addresses are automatically assigned to clusters during cluster creation basically based on the following formula: Quantity of IP addresses = Number of cluster nodes + 2 (one for Manager; one for the DB). In addition, if the Hadoop, Hue, Sqoop, and Presto or Loader and Presto components are selected during cluster deployment, one IP address is added for each component. To buy a ClickHouse cluster independently, the number of IP addresses required is calculated as follows: Number of IP addresses = Number of cluster nodes + 1 (for Manager).	-
Security Group	<p>A security group is a set of ECS access rules. It provides access policies for ECSs that have the same security protection requirements and are mutually trusted in a VPC.</p> <p>When you create an MRS cluster, a security group is automatically created by default. You can also select an existing security group from the drop-down list.</p> <p>NOTE</p> <p>When you select a security group created by yourself, ensure that the inbound rule contains a rule in which Protocol is set to All, Port is set to All, and Source is set to a trusted accessible IP address range. Do not use 0.0.0.0/0 as a source address. Otherwise, security risks may occur. If you do not know the trusted accessible IP address range, select Auto create.</p>	Auto create


Parameter	Description	Example Value
EIP	<p>After binding an EIP to an MRS cluster, you can use the EIP to access the Manager web UI of the cluster.</p> <p>When creating a cluster, you can select an available EIP from the drop-down list and bind it. If no EIP is available in the drop-down list, click Create EIPs to buy an EIP.</p> <p>NOTE The EIP must be in the same region as the cluster.</p>	Bind later

MRS Cluster Node Configuration

Table 3-4 Cluster node information

Parameter	Description	Example Value
CPU Architecture	CPU architecture type of an MRS cluster node. The value can be x86 or Kunpeng . This parameter is not available for MRS 3.1.0 and 3.1.5.	x86
Common Node Configurations	This parameter is available only when Cluster Type is set to Custom . Value options include Compact , Full-size , and OMS-separate . For details, see MRS Cluster Deployment Types .	Compact
Node Group	<p>Name of the node group in the cluster.</p> <ul style="list-style-type: none">The name of the Master node group is fixed to master_node_default_group.The system automatically creates a Core node group based on the components contained in the cluster. For example, if you select the ClickHouse component, the system adds the ClickHouse node group and deploys the ClickHouseServer role in the node group by default.If Cluster Type is set to Custom, you can customize the names of other node groups.If the data volume does not change much but the service processing capability changes greatly, manually add a Task node group. For details, see Manually Adding a Task Node Group During MRS Cluster Creation.	node_group_1

Parameter	Description	Example Value
Node Type	<p>If Cluster Type is set to Custom, you can select the node type of a non-Master node group. If the node group type is set to Task, only the NodeManager role (except the mandatory roles of nodes) can be deployed in the node group.</p>	Core
Payment Type	<p>Billing mode of nodes in a cluster.</p> <ul style="list-style-type: none"> • The billing mode of the Master and Core node groups is the same as that of the cluster. • The billing mode of the Task node group is fixed to pay-per-use. 	Pay-per-use
Node Count	<p>Configure node quantity in each node group.</p> <ul style="list-style-type: none"> • The number of nodes in a master node group ranges from 3 to 9. • There must be at least one Core node and the total number of Core and Task nodes must not exceed 10,000. • If Cluster Type is set to Custom, you can click Add Node Group to add node groups. <p>NOTE A small number of nodes may cause clusters to run slowly while a large number of nodes may be unnecessarily costly. Set an appropriate value based on data to be processed.</p>	-

Parameter	Description	Example Value
Instance Specifications	<p>Select the instance specifications of the MRS cluster node. You can click  to adjust the specifications.</p> <p>For details about the MRS cluster node specifications, see MRS Cluster Node Specifications.</p> <p>NOTE</p> <ul style="list-style-type: none"> • More advanced instance specifications provide better data processing. However, they require higher cluster cost. • Instance specifications may vary in different AZs. If no instance specifications in the current AZ can meet your requirements, switch to another AZ. • If you select HDDs for Core nodes, there is no billing information for data disks. The fees are charged with ECSs. • If you select non-HDD disks for Core nodes, the disk types of Master and Core nodes are determined by Data Disk. • If Sold out appears next to an instance specification of a node, the node of this specification cannot be bought. You can only buy nodes of other specifications. • The Master node specification (4 vCPUs 8 GB) is not within the SLA after-sales scope. It is applicable only to the test environment and is not recommended for the production environment. • For MRS 3.x or later, the memory of the master node must be greater than 64 GB. 	-
System Disk	<p>You can adjust the storage type and space of the system disk on a node as required. For details about the MRS cluster storage, see Disk Roles.</p>	-
Data Disk	<p>Storage type and space of data disks on a node. To increase the data storage capacity, you can add disks during cluster creation. A maximum of 10 disks can be added to each Core or Task node. For more information about MRS cluster storage, see Disk Roles.</p>	-
LVM	<p>This parameter is valid when a streaming Core node is created only. Click this parameter to enable or disable the disk LVM management function. This parameter is not available in MRS 3.x and later versions.</p> <p>If LVM is enabled, all disks on a node are mounted as logical volumes. This delivers more proper disk planning to avoid data skew, thereby improving system stability.</p>	Disable this function.

Parameter	Description	Example Value
Topology Adjustment	<p>If Cluster Type is set to Custom, you can adjust the deployment of each component in the cluster in the node group.</p> <p>Set Topology Adjustment to Enable and adjust the instance deployment mode based on service requirements. For details, see Role Deployment Rules for MRS Clusters.</p>	Disable this function.

Other MRS Cluster Configuration Parameters

Table 3-5 Other configuration parameters

Parameter	Description	Example Value
Kerberos Authentication	<p>Whether to enable Kerberos authentication for each component in the MRS cluster. If Kerberos authentication is enabled, users can access component resources only after being authenticated.</p> <p>This function cannot be changed after you buy a cluster.</p>	Switch this function on.
Username	Name of the administrator of Manager. admin is used by default.	admin
Password/Confirm Password	<p>Password of the Manager administrator admin. Keep the password secure.</p> <ul style="list-style-type: none"> • Must contain 8 to 26 characters. • Must contain at least four of the following: <ul style="list-style-type: none"> - Lowercase letters - Uppercase letters - Digits - At least one of the following special characters: `~!@#\$\$%^&*()-_+= [{]}:;','<.>/? • Cannot be the same as the username or the username spelled backwards. 	-

Parameter	Description	Example Value
Login Mode	<p>Method for logging in to a node in the MRS cluster.</p> <ul style="list-style-type: none"> • Password You can log in to the node as user root using a password. You need to customize the password of user root. • Key Pair Select a key pair from the drop-down list. Select "I acknowledge that I have obtained private key file <i>SSHkey-xxx</i> and that without this file I will not be able to log in to my ECS." Click View Key Pair to create or import a key pair, and then obtain the private key file. 	Password
Kerberos Encryption Type	<p>Encryption algorithm and method used by Kerberos. (This parameter is supported in MRS 3.3.1-LTS and later versions. In earlier versions, the default Kerberos encryption type is aes256-sha1,aes128-sha1.)</p> <ul style="list-style-type: none"> • aes256-sha1,aes128-sha1: indicates that the encryption algorithm and mode are AES256-CTS-HMAC-SHA1-96 AES128-CTS-HMAC-SHA1-96. • aes256-sha2,aes128-sha2: indicates that the encryption algorithm and mode are AES256-CTS-HMAC-SHA384-192 AES128-CTS-HMAC-SHA256-128. <p>NOTE If the encryption types of the two clusters are different, mutual trust cannot be set up across the Managers of clusters.</p>	aes256-sha2,aes128-sha2
Set Advanced Options	<p>Advanced function parameters of the MRS cluster. For details, see Table 3-6.</p>	-
Enterprise Project	<p>Select the enterprise project to which the cluster belongs. To use an enterprise project, create one on the Enterprise > Project Management page.</p> <p>The Enterprise Management console is designed for resource management. It helps you manage cloud-based personnel, resources, permissions, and finance in a hierarchical manner, such as management of companies, departments, and projects.</p>	default

Parameter	Description	Example Value
Required Duration	This parameter is valid in Yearly/Monthly billing mode and indicates a cluster subscription duration. The minimum cluster duration is 1 month and the maximum available cluster duration is 1 year. If Auto-renew is selected, monthly subscriptions are automatically renewed every month and yearly subscriptions are automatically renewed every year.	-
Secure Communications	To allow the MRS console to access big data components in the user VPC, you need to activate the relevant security group rules. For details, see Configuring Secure Communication Authorization for an MRS Cluster .	Select this function.

Table 3-6 MRS cluster advanced configuration

Parameter	Description	Example Value
Hostname Prefix	Enter the prefix for the computer hostname of an ECS or BMS in the cluster.	-
Cryptographic Algorithm	Algorithm used for encrypting and decrypting passwords in the cluster system. <ul style="list-style-type: none"> • International: general encryption algorithm • Chinese: SM series cryptographic algorithms are compatible with general cryptographic algorithms. 	International
Tag	It is recommended that you use the tag function of TMS to add the same tag to different cloud resources.	-
Auto Scaling	Configure an auto scaling policy for the Task node group after the cluster is created. For details, see MRS Task Node Auto Scaling .	-
Bootstrap Action	Bootstrap actions run scripts on specified nodes in a cluster to install third-party software and modify the cluster's running environment. For details, see Adding MRS Node Bootstrap Actions and Installing Third-Party Software .	-

Parameter	Description	Example Value
Agency	<p>When an agency is bound to an ECS or BMS, it can manage designated resources in the cluster. It is essential to consider service requirements before deciding to configure an agency.</p> <p>For example, you can configure an ECS agency to automatically obtain the AK/SK to access OBS. For details, see Interconnecting an MRS Cluster with OBS Using an IAM Agency.</p> <p>To bind an agency to an MRS cluster, you need to create an IAM agency with required permissions in advance. By default, the system generates an MRS_ECS_DEFAULT_AGENCY agency. This agency has the OBSOperateAccess permission and the CESFullAccess (for users who have enabled fine-grained policies), CES Administrator, and KMS Administrator permissions in the region where the cluster is located.</p>	Bind later
Metric Sharing	Monitoring metrics of big data components are collected. If a fault occurs when you use a cluster, share the monitoring metrics with Huawei Cloud technical support for troubleshooting.	Disable

Parameter	Description	Example Value
System Disk Encryption	<p>Whether to encrypt data in the system disk mounted to the MRS cluster node. This function is disabled by default.</p> <p>Keys used by encrypted system disks are provided by Key Management Service (KMS) in Data Encryption Workshop (DEW). You do not need to build and maintain the key management infrastructure. To enable this function, you must have the Security Administrator and KMS Administrator permissions.</p> <p>For more information about disk encryption, see Managing Encrypted EVS Disks.</p> <p>Configure the following parameters to enable this function:</p> <ul style="list-style-type: none">• Data Disk Key ID: key ID of the selected key name.• Data Disk Key Name: Select the name of the key used to encrypt the data disk. By default, the default master key named evs/default is selected. You can select another master key from the drop-down list. <p>If cloud disks are encrypted using a master key and it is then disabled or scheduled for deletion, the cloud disks can no longer be read from or written to, and data on these disks may never be restored. Exercise caution when performing this operation.</p> <p>Click View the Key List to enter a page where you can create and manage keys.</p>	Disable

Parameter	Description	Example Value
Data Disk Encryption	<p>Whether to encrypt data in the data disk mounted to the MRS cluster node. This function is disabled by default. MRS 3.1.0 and MRS 3.1.2-LTS.3 do not support this function.</p> <p>For more information about disk encryption, see Managing Encrypted EVS Disks.</p> <p>Keys used by encrypted system disks are provided by Key Management Service (KMS) in Data Encryption Workshop (DEW). You do not need to build and maintain the key management infrastructure. To enable this function, you must have the Security Administrator and KMS Administrator permissions.</p> <p>Configure the following parameters to enable this function:</p> <ul style="list-style-type: none"> • Data Disk Key ID: key ID of the selected key name. • Data Disk Key Name: Select the name of the key used to encrypt the data disk. By default, the default master key named evs/default is selected. You can select another master key from the drop-down list. <p>If cloud disks are encrypted using a master key and it is then disabled or scheduled for deletion, the cloud disks can no longer be read from or written to, and data on these disks may never be restored. Exercise caution when performing this operation.</p> <p>Click View the Key List to enter a page where you can create and manage keys.</p>	Disable
Alarm	<p>If the alarm function is enabled, the cluster maintenance personnel can be notified in a timely manner to locate faults when the cluster runs abnormally or the system is faulty. To send alarm messages, you need to enable the Simple Message Notification (SMN) service, set a notification rule, and bind the rule to an SMN topic.</p> <ul style="list-style-type: none"> • Rule Name: Name of the rule for sending alarm messages. The value can contain only digits, letters, hyphens (-), and underscores (_). • Topic Name: Select an existing SMN topic or click Create Topic to create a topic. To deliver messages published to a topic, you need to add a subscriber to the topic. For details, see Adding Subscriptions to a Topic. 	Enable

Parameter	Description	Example Value
Install UniAgent	Install UniAgent on MRS cluster nodes to simplify plug-in management, provide AOM instructions, and enable script delivery and execution.	No
Logging	Whether to collect logs when cluster creation fails. After the logging function is enabled, system logs and component run logs are automatically collected and saved to the OBS file system in scenarios such as cluster creation failures and scale-out or scale-in failures for O&M personnel to quickly locate faults. The log information is retained for a maximum of seven days.	Disable

Manually Adding a Task Node Group During MRS Cluster Creation

To add a Task node group in the current MRS cluster for auto scaling, do as follows:

Add an analysis Task node group.

- For a cluster whose **Cluster Type** is **Analysis cluster** or **Hybrid cluster**:
The system automatically adds the analysis Task node group **task_node_analysis_group**. Set the number of nodes, instance specifications, and node disk configurations as needed.
If the analysis Task node group is not required, you can delete it.
- For a cluster whose **Cluster Type** is **Custom**:
 - a. In the cluster node configuration area, click **Add Node Group** to manually add a node group.
 - b. Set **Topology Adjustment** to **Enable**.
 - c. Select the role topology so that the node group contains only the NodeManager (NM) role.

Add a Task node group after the cluster is created. For details, see [Adding a Task Node](#).

Adding streaming Task node group

For a streaming or hybrid cluster that contains the Storm component:

The system automatically adds the streaming Task node group **task_node_streaming_group**. Set the number of nodes, instance specifications, and node disk configurations as needed.

If the streaming Task node group is not required, you can manually delete it.

Viewing Failed MRS Cluster Creation Tasks

If a cluster fails to be created, the failed task will be managed on the **Manage Failed Tasks** page.


On the cluster list page, click  to switch to the **Manage Failed Tasks** page. In the **Task Status** column, hover the cursor over the task status to view the failure cause.

Table 3-7 lists the error codes of MRS cluster creation failures.

Table 3-7 Error codes

Error Code	Description
MRS.101	Insufficient quota to meet your request. Contact customer service to increase the quota.
MRS.102	The token cannot be null or invalid. Try again later or contact customer service.
MRS.103	Invalid request. Try again later or contact customer service.
MRS.104	Insufficient resources. Try again later or contact customer service.
MRS.105	Insufficient IP addresses in the existing subnet. Try again later or contact customer service.
MRS.201	Failed due to an ECS error. Try again later or contact customer service.
MRS.202	Failed due to an IAM error. Try again later or contact customer service.
MRS.203	Failed due to a VPC error. Try again later or contact customer service.
MRS.400	MRS system error. Try again later or contact customer service.

4 Installing an MRS Cluster Client

4.1 Installing a Client (MRS 3.x)

You need to install a cluster client to connect to the component server in the cluster and perform tasks such as component connection and job submission. You can install the cluster client on a node within the cluster or on a node outside of it.

To ensure that certain features function properly, it is important to reinstall the client after modifying the server configuration of a component in the cluster. This ensures that the client version matches the server version.

This section describes how to install the MRS 3.x cluster client. For MRS 2.x and earlier versions, see [Installing a Client \(MRS 2.x or Earlier\)](#). If the cluster contains the Flume component, install the Flume client separately. For details, see [Installing the Flume Client](#).

Prerequisites

- If the node where the client is to be installed is outside the cluster, the node must be able to communicate with the nodes in the MRS cluster. Otherwise, client installation will fail.
- The node where the client is to be installed must have the NTP service enabled and synchronized time with the MRS cluster server. Otherwise, client installation will fail.
- You can install the client on the node as a **root** or any OS user. The user must have the operation permission on the client file storage directory and installation directory.
- When you install the client as a user other than **omm** or **root**, and the **/var/tmp/patch** directory already exists, you have changed the permission for the directory to **777** and changed the permission for the logs in the directory to **666**.

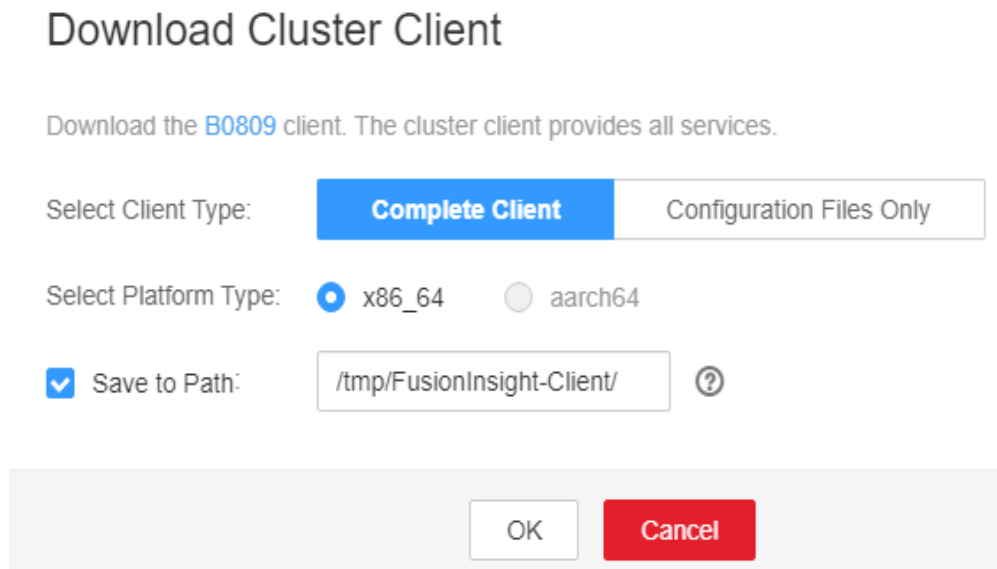
Installing a Client on a Node in a Cluster

Step 1 Obtain the client software package.

Log in to Manager by referring to [Accessing MRS Manager](#). Click **Cluster**. On the **Dashboard** page, click the more sign (...) and select **Download Client**. In the

displayed **Download Cluster Client** dialog box, configure parameters and click **OK**.

Figure 4-1 Downloading a client



NOTE

- The client software package downloaded from the FusionInsight Manager homepage contains the clients of all services (excluding Flume) in the cluster. To download the client of a single service, choose **Cluster > Services > Service name**, click **More**, and select **Download Client**.
- For MRS 3.3.0 or later, click **Download Client** on the Manager homepage.

Table 4-1 Client download parameters

Parameter	Description	Example Value
Select Client Type	<ul style="list-style-type: none"> • Complete Client: contains the complete client software package and configuration files. Generally, this option is selected. • Configuration Files Only: downloads only client configuration files in the scenario where the administrator modifies the component server configuration on FusionInsight Manager after the complete client is downloaded and installed in an application development task, and developers need to update client configuration files. 	Complete Client

Parameter	Description	Example Value
Select Platform Type	<p>The client type must match the architecture of the node where the client is to be installed. Otherwise, the installation fails.</p> <p>For clusters of the LTS version, only the client software package whose type is the same as that of FusionInsight Manager can be downloaded.</p> <ul style="list-style-type: none">• x86_64: indicates the client software package that can be deployed on a x86 platform.• aarch64: indicates the client software package that can be deployed on a Kunpeng server.	x86_64
Save to Path	<p>The path for storing the client software package on the active OMS node</p> <ul style="list-style-type: none">• Select Save to Path: Customize the path for storing the client software package on the active OMS node. User omm must have the read, write, and execute permissions on the path. If the path is not changed, the client file generated is saved in the /tmp/FusionInsight-Client directory on the active OMS node in the cluster by default.• Not to select Save to Path: The generated client file is automatically downloaded and saved to the local host. Before installing the client, you need to upload the file to a specified directory on the target node.	Select Save to Path

Step 2 Copy the client software package to a specified directory on the node where the client is to be installed.

By default, the client software package is stored on the active OMS node in the cluster. To install the client on other nodes in the cluster, log in to the active OMS node as user **omm** and run the following command to copy the software package to the specified node. Otherwise, skip this step.

For example, copy the software package to the **/tmp/clienttemp** directory.

```
scp -p /tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_Client.tar IP  
address of the node where the client is to be installed:/tmp/clienttemp
```

Step 3 Log in to the node where the client is to be installed as the client installation user.

 **NOTE**

You can install the client on the node as user **root**, **omm**, or any other OS user. The user needs to have the operation permission on the directory for storing the client file and the installation directory. The permission on the two directories is **755**.

Step 4 Decompress the client software package.

1. Go to the directory where the package is stored, for example, `/tmp/clienttemp`.
cd /tmp/clienttemp
2. Decompress the installation package.
tar -xvf FusionInsight_Cluster_1_Services_Client.tar
3. Run the **sha256sum** command to verify the decompressed file and check whether the command output is consistent with the content in the **sha256** file.
sha256sum -c FusionInsight_Cluster_1_Services_ClientConfig.tar.sha256
FusionInsight_Cluster_1_Services_ClientConfig.tar: OK
4. Decompress the installation package.
tar -xvf FusionInsight_Cluster_1_Services_ClientConfig.tar

Step 5 Go to the directory where the client software package is decompressed and run the following command to install the client to the specified directory:**cd FusionInsight_Cluster_1_Services_ClientConfig****./install.sh** *Client installation directory*

Example:

./install.sh /opt/hadoopclient

Wait until the client installation is complete.

...
The component client is installed successfully**NOTE**

- **When installing the client, you can choose the installation directory. If you opt to use an existing directory, it must be empty. Also, the installation directory cannot have spaces and should only contain uppercase and lowercase letters, digits, and underscores (_).**
- Delete the client installation directory when uninstalling a client.
- To ensure that the client can be used only by the current user, add the **-o** parameter during the installation. For example, run the **./install.sh /opt/hadoopclient -o** command to install the client.

Step 6 Check whether the client is installed.

1. Go to the client installation directory and load environment variables.
cd /opt/hadoopclient
source bigdata_env
2. Run related commands based on the cluster mode.
 - If Kerberos authentication is not enabled for the cluster, you can directly run commands related to the component client.
View files in the HDFS root directory:
hdfs dfs -ls /
 - If Kerberos authentication is enabled for the cluster, run the **kinit** command to perform user authentication.

Example:

kinit admin

Password for xxx@HADOOP.COM: #Enter the password of user **admin**.

Run the **klist** command to query and confirm authentication details.

Ticket cache: FILE:/tmp/krb5cc_0
Default principal: xxx@HADOOP.COM

Valid starting	Expires	Service principal
...		

----End

Installing a Client on a Node Outside a Cluster

Step 1 Prepare a Linux ECS for installing the MRS cluster client.

- The recommended ECS OSs and versions are as follows:

Table 4-2 OS reference list


CPU Architecture	OS	Supported Version
x86 computing	EulerOS	EulerOS 2.5
	SUSE	SUSE Linux Enterprise Server 12 SP4 (SUSE 12.4)
	Red Hat	Red Hat-7.5-x86_64 (Red Hat 7.5)
	CentOS	CentOS 7.6
Kunpeng computing	EulerOS	EulerOS 2.8
	CentOS	CentOS 7.6

- Sufficient disk space (at least 40 GB) must be allocated to the ECS client installation directory.
- The ECS must be in the same VPC and security group as the MRS cluster.
- All ports in the inbound direction of the MRS cluster security group are open to the client node. For details, see [Adding a Security Group Rule](#).
- The NTP service has been installed on the ECS OS and is running properly. If the NTP service is not installed, run the **yum install ntp -y** command to install it when the **yum** source is configured.
- The ECS must allow users to log in to it using a password (SSH).

Step 2 Configure NTP time synchronization for the node where the client is to be installed to synchronize time with the MRS cluster.

- Log in to the MRS console and click the MRS cluster name from the cluster list.
- Click **Nodes** to expand the Master node group list and view the IP addresses of the **Master1** and **Master2** nodes in the cluster.

Figure 4-2 Viewing the IP address of the Master node



The screenshot shows a web interface for a cluster named 'master_node_default_group'. It displays a table with columns for 'Node Name/Resource ID', 'IP', and 'Operating Status'. Two nodes are listed: 'node-master1BBSg' with IP '192.168.12.136' and 'node-master2gUSI' with IP '192.168.12.61'. Both nodes are in a 'Running' state.


Node Name/Resource ID	IP	Operating Status
★ node-master1BBSg	192.168.12.136	Running
☆ node-master2gUSI	192.168.12.61	Running

3. Log in to the node where the client is to be installed as user **root** and run the following command to edit the NTP configuration file. Add the IP addresses of the **Master1** and **Master2** nodes in the MRS cluster and comment out other **server** addresses.

```
vi /etc/ntp.conf
```

```
server master1_ip prefer
server master2_ip
```

Figure 4-3 Modifying the NTP configuration file



The screenshot shows the contents of the /etc/ntp.conf file. The configuration includes comments and settings for driftfile, restrict, and server. The following lines are highlighted with a red box:

```
#server 0.centos.pool.ntp.org iburst
#server 1.centos.pool.ntp.org iburst
#server 2.centos.pool.ntp.org iburst
#server 3.centos.pool.ntp.org iburst
server 10.9.2.38 prefer
server 10.9.2.39
```

4. Save the configuration file and run the following command to disable the NTP service:

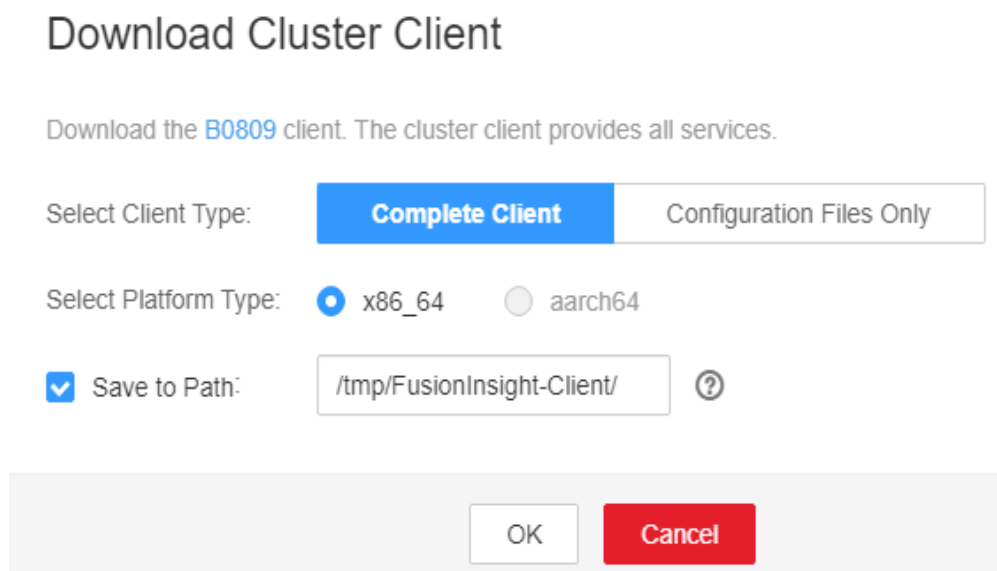
```
service ntpd stop
```

5. Manually synchronize the time.
`/usr/sbin/ntpdate IP address of the active Master node`
6. Start the NTP service.
`service ntpd start`
Or
`systemctl restart ntpd`
7. Run the `ntpstat` command to check the time synchronization result.
synchronised to NTP server (xxx) at stratum 2
time correct to within 12 ms
polling server every 16 s

Step 3 Obtain the client software package.

Log in to Manager by referring to [Accessing MRS Manager](#). Click **Cluster**. On the **Dashboard** page, click the more sign (...) and select **Download Client**. In the displayed **Download Cluster Client** dialog box, configure parameters and click **OK**.

Figure 4-4 Downloading a client



NOTE

- The client software package downloaded from the FusionInsight Manager homepage contains the clients of all services (excluding Flume) in the cluster. To download the client of a single service, choose **Cluster > Services > Service name**, click **More**, and select **Download Client**.
- For MRS 3.3.0 or later, click **Download Client** on the Manager homepage.

Table 4-3 Client download parameters

Parameter	Description	Example Value
Select Client Type	<ul style="list-style-type: none">• Complete Client: contains the complete client software package and configuration files. Generally, this option is selected.• Configuration Files Only: downloads only client configuration files in the scenario where the administrator modifies the component server configuration on FusionInsight Manager after the complete client is downloaded and installed in an application development task, and developers need to update client configuration files.	Complete Client
Select Platform Type	<p>The client type must match the architecture of the node where the client is to be installed. Otherwise, the installation fails.</p> <p>For clusters of the LTS version, only the client software package whose type is the same as that of FusionInsight Manager can be downloaded.</p> <ul style="list-style-type: none">• x86_64: indicates the client software package that can be deployed on a x86 platform.• aarch64: indicates the client software package that can be deployed on a Kunpeng server.	x86_64
Save to Path	<p>The path for storing the client software package on the active OMS node</p> <ul style="list-style-type: none">• Select Save to Path: Customize the path for storing the client software package on the active OMS node. User omm must have the read, write, and execute permissions on the path. If the path is not changed, the client file generated is saved in the /tmp/FusionInsight-Client directory on the active OMS node in the cluster by default.• Not to select Save to Path: The generated client file is automatically downloaded and saved to the local host. Before installing the client, you need to upload the file to a specified directory on the target node.	Select Save to Path

Step 4 Copy the client software package to a specified directory on the node where the client is to be installed.

By default, the client software package is stored on the active OMS node in the cluster. To install the client on other nodes in the cluster, log in to the active OMS node as user **omm** and run the following command to copy the software package to the specified node. Otherwise, skip this step.

For example, copy the software package to the **/tmp/clienttemp** directory.

```
scp -p /tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_Client.tar IP  
address of the node where the client is to be installed:/tmp/clienttemp
```

Step 5 Log in to the node where the client is to be installed as the client installation user.

 **NOTE**

You can install the client on the node as user **root**, **omm**, or any other OS user. The user needs to have the operation permission on the directory for storing the client file and the installation directory. The permission on the two directories is **755**.

Step 6 Decompress the client software package.

1. Go to the directory where the package is stored, for example, **/tmp/clienttemp**.

```
cd /tmp/clienttemp
```

2. Decompress the installation package.

```
tar -xvf FusionInsight_Cluster_1_Services_Client.tar
```

3. Run the **sha256sum** command to verify the decompressed file and check whether the command output is consistent with the content in the **sha256** file.

```
sha256sum -c FusionInsight_Cluster_1_Services_ClientConfig.tar.sha256
```

```
FusionInsight_Cluster_1_Services_ClientConfig.tar: OK
```

4. Decompress the installation package.

```
tar -xvf FusionInsight_Cluster_1_Services_ClientConfig.tar
```

Step 7 Check the network connection of the client.

1. To install the client on a node, map the host names and IP addresses of all nodes in the cluster in the **/etc/hosts** file. Here are the steps to import the domain name mapping of the cluster to the **hosts** file:

- a. Switch to user **root** or a user who has the permission to modify the **hosts** file.

```
su - root
```

- b. Go to the directory where the client package is decompressed.

```
cd /tmp/clienttemp/FusionInsight_Cluster_1_Services_ClientConfig
```

- c. Import the domain name mapping to the **hosts** file.

```
cat realm.ini >> /etc/hosts
```


 NOTE

- If the host where the client is installed is not a node in the cluster, configure network connections for the client to prevent errors when you run commands on the client.
- If Spark tasks are executed in yarn-client mode, add the **spark.driver.host** parameter to the file *Client installation directory/Spark/spark/conf/spark-defaults.conf* and set the parameter to the client IP address.
- If the yarn-client mode is used, you need to configure the mapping between the IP address and host name of the client in the **hosts** file on the active and standby YARN nodes (ResourceManager nodes in the cluster) to make sure that the Spark web UI is properly displayed.

Step 8 Go to the directory where the client software package is decompressed and run the following command to install the client to the specified directory:

```
cd FusionInsight_Cluster_1_Services_ClientConfig
```

```
./install.sh Client installation directory
```

Example:

```
./install.sh /opt/hadoopclient
```

Wait until the client installation is complete.

...

```
The component client is installed successfully
```

 NOTE

- **When installing the client, you can choose the installation directory. If you opt to use an existing directory, it must be empty. Also, the installation directory cannot have spaces and should only contain uppercase and lowercase letters, digits, and underscores (_).**
- Delete the client installation directory when uninstalling a client.
- To ensure that the client can be used only by the current user, add the **-o** parameter during the installation. For example, run the **./install.sh /opt/hadoopclient -o** command to install the client.

Step 9 Check whether the client is installed.

1. Go to the client installation directory and load environment variables.

```
cd /opt/hadoopclient
```

```
source bigdata_env
```

2. Run related commands based on the cluster mode.
 - If Kerberos authentication is not enabled for the cluster, you can directly run commands related to the component client.

View files in the HDFS root directory:

```
hdfs dfs -ls /
```

- If Kerberos authentication is enabled for the cluster, run the **kinit** command to perform user authentication.

Example:

```
kinit admin
```

```
Password for xxx@HADOOP.COM: #Enter the password of user admin.
```

Run the **klist** command to query and confirm authentication details.

```
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: xxx@HADOOP.COM

Valid starting    Expires          Service principal
...

```

----End

4.2 Installing a Client (MRS 2.x or Earlier)

You need to install a cluster client to connect to the component server in the cluster and perform tasks such as component connection and job submission. You can install the cluster client on a node within the cluster or on a node outside of it.

To ensure that certain features function properly, it is important to reinstall the client after modifying the server configuration of a component in the cluster. This ensures that the client version matches the server version.

This section describes how to install the cluster client of MRS 2.x or earlier. For details about MRS 3.x, see [Installing a Client \(MRS 3.x\)](#).

After an MRS cluster earlier than 3.x is created, the client is installed on the active Master node by default. The installation directory is `/opt/client`.

Installing a Client on a Node in a Cluster

Step 1 Log in to Manager by referring to [Accessing MRS Manager](#), choose **Services** > **Download Client**, and download the client installation package to the active OMS node.

Step 2 On the MRS console, view the IP address on the **Nodes** page of the specified cluster.

Record the IP address of the node where the client is to be installed and the IP address of the active Master node (active OMS node).

Step 3 Log in to the active OMS node as user **root** and run the following command to switch to user **omm**.

```
sudo su - omm
```

Step 4 Copy the client installation package to the specified node.

```
scp -p /tmp/MRS-client/MRS_Services_Client.tar IP address of the node where
the client is to be installed:/opt/client_tmp
```

Step 5 Log in to the node where the client is to be installed as the **root** user.

Step 6 Install the client.

```
cd /opt/client_tmp
tar -xvf MRS_Services_Client.tar
tar -xvf MRS_Services_ClientConfig.tar
cd MRS_Services_ClientConfig
./install.sh Client installation directory
```

Example:

./install.sh /opt/hadoopclient

Wait until the client is installed.

 **NOTE**

- **When installing the client, you can choose the installation directory. If you opt to use an existing directory, it must be empty. Also, the installation directory cannot have spaces and should only contain uppercase and lowercase letters, digits, and underscores (_).**
- Delete the client installation directory when uninstalling a client.

Step 7 Check whether the client is installed.

1. Go to the client installation directory and load environment variables.

cd /opt/hadoopclient

source bigdata_env

2. Run related commands based on the cluster mode.
 - If Kerberos authentication is not enabled for the cluster, you can directly run commands related to the component client.

View files in the HDFS root directory:

hdfs dfs -ls /

- If Kerberos authentication is enabled for the cluster, run the **kinit** command to perform user authentication.

Example:

kinit admin

Password for xxx@HADOOP.COM: #Enter the password of user **admin**.

Run the **klist** command to query and confirm authentication details.

Ticket cache: FILE:/tmp/krb5cc_0
Default principal: xxx@HADOOP.COM

Valid starting	Expires	Service principal
...		

----End

Installing a Client on a Node Outside the Cluster

Step 1 Prepare a Linux ECS for installing the MRS cluster client.

- The recommended ECS OSs and versions are as follows:

Table 4-4 OS reference list

CPU Architecture	OS	Supported Version
x86 computing	EulerOS	<ul style="list-style-type: none"> - Available: EulerOS 2.2 - Available: EulerOS 2.3 - Available: EulerOS 2.5

CPU Architecture	OS	Supported Version
Kunpeng computing (Arm)	EulerOS	Available: EulerOS 2.8

- The CPU architecture of the ECS must be the same as that of the MRS cluster node.
- Sufficient disk space (at least 40 GB) must be allocated to the ECS client installation directory.
- The ECS must be in the same VPC and security group as the MRS cluster.
- All ports in the inbound direction of the MRS cluster security group are open to the client node. For details, see [Adding a Security Group Rule](#).
- The NTP service has been installed on the ECS OS and is running properly. If the NTP service is not installed, run the **yum install ntp -y** command to install it when the **yum** source is configured.
- The ECS must allow users to log in to it using a password (SSH).

Step 2 Log in to Manager of the cluster by referring to [Accessing MRS Manager](#) and choose **Services > Download Client**.

- In **Client Type**, select **All client files**.
- In **Download to**, select **Remote host**.
- Set **Host IP Address** to the IP address of the node where the client is to be installed, **Host Port** to **22**, and **Save Path** to **/tmp**.
If the default SSH login port of the node has been changed, set **Host Port** to the new port.
- Set **Login User** to **root**.
If other users are used, ensure that the users have read, write, and execute permission on the save path.
- Select **Password** or **SSH Private Key** for **Login Mode**.
 - **Password**: Enter the password of user **root** set during cluster creation.
 - **SSH Private Key**: Select and upload the key file used for creating the cluster.

Step 3 Click **OK** to generate a client file.

After the file download is confirmed, obtain the client software package from the specified path on the node where the client will be installed.

If the client download fails, check the username, password, and security group configuration of the remote host. Ensure that the username and password are accurate and that the inbound rule for the SSH port has been added to the remote host's security group. Once confirmed, retry the client download.

 NOTE

Generating a client will occupy a large number of disk I/Os. You are advised not to download a client when the cluster is being installed, started, and patched, or in other unstable states.

- Step 4** Configure NTP time synchronization for the node where the client is to be installed to synchronize time with the MRS cluster.
1. Log in to the MRS console and click the MRS cluster name from the cluster list.
 2. Click **Nodes** to expand the Master node group list and view the IP addresses of the **Master1** and **Master2** nodes in the cluster.
 3. Log in to the node where the client is to be installed as user **root** and run the following command to edit the NTP configuration file. Add the IP addresses of the **Master1** and **Master2** nodes in the MRS cluster and comment out other **server** addresses.

```
vi /etc/ntp.conf
```

```
server master1_ip prefer
server master2_ip
```

Figure 4-5 Modifying the NTP configuration file

```
# For more information about this file, see the man pages
# ntp.conf(5), ntp_acc(5), ntp_auth(5), ntp_clock(5), ntp_misc(5), ntp_mon(5).

driftfile /var/lib/ntp/drift

# Permit time synchronization with our time source, but do not
# permit the source to query or modify the service on this system.
restrict default nomodify notrap nopeer noquery

# Permit all access over the loopback interface. This could
# be tightened as well, but to do so would effect some of
# the administrative functions.
restrict 127.0.0.1
restrict ::1

# Hosts on local network are less restricted.
#restrict 192.168.1.0 mask 255.255.255.0 nomodify notrap

# Use public servers from the pool.ntp.org project.
# Please consider joining the pool (http://www.pool.ntp.org/join.html).
#server 0.centos.pool.ntp.org iburst
#server 1.centos.pool.ntp.org iburst
#server 2.centos.pool.ntp.org iburst
#server 3.centos.pool.ntp.org iburst
#server 4.centos.pool.ntp.org iburst
server 10.9.2.38 prefer
server 10.9.2.39
#broadcast 192.168.1.255 autokey           # broadcast server
#broadcastclient                          # broadcast client
#broadcast [redacted] autokey             # multicast server
#multicastclient [redacted]              # multicast client
#manycastserver [redacted]               # manycast server
#manycastclient [redacted] autokey       # manycast client

# Enable public key cryptography.
#crypto
```

4. Save the configuration file and run the following command to disable the NTP service:
service ntpd stop
5. Manually synchronize the time.
/usr/sbin/ntpdate IP address of the active Master node

6. Start the NTP service.
service ntpd start
Or
systemctl restart ntpd
7. Run the **ntpstat** command to check the time synchronization result.
synchronised to NTP server (xxx) at stratum 2
time correct to within 12 ms
polling server every 16 s

Step 5 Copy the downloaded client software package to the **/opt** directory.

```
cp /tmp/MRS_Services_Client.tar /opt
```

Step 6 Decompress the software package.

```
cd /opt  
tar -xvf MRS_Services_Client.tar
```

Step 7 Verify the file.

```
sha256sum -c MRS_Services_ClientConfig.tar.sha256
```

The command output is as follows:

```
MRS_Services_ClientConfig.tar: OK
```

Step 8 Decompress the software package.

```
tar -xvf MRS_Services_ClientConfig.tar
```

Step 9 Install the cluster client.

```
sh /opt/MRS_Services_ClientConfig/install.sh Client installation directory
```

Example:

```
sh /opt/MRS_Services_ClientConfig/install.sh /opt/hadoopclient
```

If the following information is displayed, the client has been successfully installed:

```
Components client installation is complete.
```

 **NOTE**

- When installing the client, you can choose the installation directory. If you opt to use an existing directory, it must be empty. Also, the installation directory cannot have spaces and should only contain uppercase and lowercase letters, digits, and underscores (_).
- Delete the client installation directory when uninstalling a client.

Step 10 Check whether the client is installed.

1. Go to the client installation directory and load environment variables.

```
cd /opt/hadoopclient  
source bigdata_env
```

2. Run related commands based on the cluster mode.

- If Kerberos authentication is not enabled for the cluster, you can directly run commands related to the component client.

View files in the HDFS root directory:

hdfs dfs -ls /

- If Kerberos authentication is enabled for the cluster, run the **kinit** command to perform user authentication.

Example:

kinit admin

Password for xxx@HADOOP.COM: #Enter the password of user **admin**.

Run the **klist** command to query and confirm authentication details.

Ticket cache: FILE:/tmp/krb5cc_0

Default principal: xxx@HADOOP.COM

Valid starting	Expires	Service principal
...		

...

----End

5 Submitting an MRS Job

5.1 MRS Job Types

Category

An MRS job is the program execution platform of MRS. It is used to process and analyze user data. You can create jobs online using the MRS console or submit jobs in the background through the cluster client.

MRS jobs typically process data from OBS or HDFS. To create a job, you must first upload the data to be analyzed to OBS. MRS utilizes the data stored in OBS for computing and analysis.

MRS allows exporting data from OBS to HDFS for computing and analyzing. After data analysis and computing is complete, you can store the data in the HDFS or export it to OBS. HDFS and OBS can store compressed data in **bz2** and **gz** formats.

You can create the following types of jobs online in an MRS cluster:

- MapReduce can quickly process large-scale data in parallel. It is a distributed data processing model and execution environment. MRS supports the submission of MapReduce JAR programs.
- Spark is a distributed in-memory computing framework. MRS supports SparkSubmit, Spark Script, and Spark SQL jobs.
 - SparkSubmit: You can submit Spark JAR and Spark Python programs, execute the Spark Application, and compute and process user data.
 - SparkScript: You can submit SparkScript scripts and batch execute Spark SQL statements.
 - Spark SQL: You can use Spark SQL statements (similar to SQL statements) to query and analyze user data in real time.
- Hive is an open-source data warehouse based on Hadoop. MRS allows you to submit HiveScript scripts and directly execute Hive SQL statements.
- Flink is a distributed big data processing engine that can perform stateful computations over both unbounded and bounded data streams.
- HadoopStreaming works similarly to a standard Hadoop job, where you can define the input and output HDFS paths, as well as the mapper and reducer executable programs.

Job Execution Permission Description

For a security cluster with Kerberos authentication enabled, a user needs to synchronize an IAM user before submitting a job on the MRS web UI. After the synchronization is completed, the MRS system generates a user with the same IAM username. Whether a user has the permission to submit jobs depends on the IAM policy bound to the user during IAM synchronization. For details about the job submission policy, see [Table 6-58](#) in [Synchronizing IAM Users to MRS](#).

When a user submits a job that involves the resource usage of a specific component, such as accessing HDFS directories and Hive tables, user **admin** (Manager administrator) must grant the relevant permission to the user.

- Step 1** Log in to Manager of the cluster as user **admin**.
- Step 2** Add the role of the component whose permission is required by the user. For details, see [Managing MRS Cluster Roles](#).
- Step 3** Change the user group to which the user who submits the job belongs and add the new component role to the user group. For details, see [Managing MRS Cluster User Groups](#).

NOTE

After the component role bound to the user group to which the user belongs is modified, it takes some time for the role permissions to take effect.

----End

5.2 Uploading Application Data to an MRS Cluster

MRS clusters generally process data from the OBS file system or the HDFS file system in the cluster. OBS provides you with the data storage capabilities that are massive, secure, reliable, and cost-effective.

You can access, manage, and use OBS data on the MRS console and OBS client. You can also import OBS data to the HDFS system of a cluster for processing. Note that the file upload rate may decrease as the file size increases and this method more suitable for scenarios with smaller amounts of data.

Importing OBS Data to HDFS

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page displayed by default, click the name of the target cluster to enter the cluster details page.

Complete IAM user synchronization first for MRS clusters with Kerberos authentication enabled. (On the **Dashboard** page of the cluster details page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- Step 3** Click **Files** to go to the file management page.
- Step 4** Select **HDFS File List**.

Figure 5-1 HDFS file list

[HDFS File List](#)

You can view HDFS audit logs on the tenant plane.

/

File Name	File Size
app-logs	--
apps	--
datasets	--
datastore	--
flume	--

Step 5 Go to the directory where the data to be imported is stored.

Click **Create** to create a folder directory or select an existing folder in HDFS.

Step 6 Click **Import Data** and configure the HDFS and OBS paths correctly.

When configuring the OBS or HDFS path, click **Browse**, select a file directory, and click **OK**.

Figure 5-2 Importing data

Import Data from OBS to HDFS

OBS Path

HDFS Path

- OBS path description:
 - The path must start with **obs://**.
 - Files or programs encrypted by KMS cannot be imported.
 - An empty folder cannot be imported.
 - The directory and file name can contain letters, digits, hyphens (-), and underscores (_), but cannot contain special characters ;|&>,<'\$*?\
 - The directory and file name cannot start or end with a space, but can contain spaces between them.
 - The OBS full path contains a maximum of 255 characters.
- HDFS path description:

- The directory and file name can contain letters, digits, hyphens (-), and underscores (_), but cannot contain the following special characters: ; | & > , < ' \$ * ? \ :
- The directory and file name cannot start or end with a space, but can contain spaces between them.
- The HDFS full path contains a maximum of 255 characters.

Step 7 Click **OK**.

You can view the file upload progress on the **File Operation Records** page. The system generates a DistCp job for processing the data import operation. You can also view the job execution status on the **Job Management** page.

----End

Exporting HDFS Data to OBS

Step 1 Log in to the MRS console.

Step 2 On the **Active Clusters** page displayed by default, click the name of the target cluster to enter the cluster details page.

Step 3 Click the **Files** tab to go to the file management page.

Step 4 Select **HDFS File List**.

Step 5 Go to the data storage directory.

Step 6 Click **Export Data** and configure the OBS and HDFS paths. When configuring the OBS or HDFS path, click **Browse**, select a file directory, and click **OK**.

Figure 5-3 Exporting data

Export Data from HDFS to OBS

The screenshot shows a dialog box titled "Export Data from HDFS to OBS". It contains two rows of input fields. The first row is labeled "HDFS Path" and has a text input field with a small icon on the left and a "Browse" button on the right. The second row is labeled "OBS Path" and has a text input field with a question mark icon on the left and a "Browse" button on the right. At the bottom of the dialog, there are two buttons: a red "OK" button and a grey "Cancel" button.

NOTE

When a folder is exported to OBS, a label file named **folder name_ \$folder\$** is added to the OBS path. Ensure that the exported folder is not empty. If the exported folder is empty, OBS cannot display the folder and only generates a file named **folder name_ \$folder\$**.

Step 7 Click **OK**.

You can view the file upload progress on the **File Operation Records** page. The system generates a DistCp job for processing the data import operation. You can also view the job execution status on the **Job Management** page.

----End

5.3 Running an MRS Job

5.3.1 Running a MapReduce Job

MRS allows you to submit and run your own programs, and get the results. This section will show you how to submit a MapReduce job in an MRS cluster.

MapReduce jobs are used to submit Hadoop JAR programs to quickly process a large amount of data in parallel. MapReduce is a distributed data processing mode.

You can create a job online and submit it for running on the MRS console, or submit a job in CLI mode on the MRS cluster client.

Prerequisites

- You have uploaded the program packages and data files required by jobs to OBS or HDFS.
- If the job program needs to read and analyze data in the OBS file system, you need to configure storage-compute decoupling for the MRS cluster. For details, see [Configuring Storage-Compute Decoupling for an MRS Cluster](#).

Submitting a Job on the Console

Step 1 Log in to the MRS console.

Step 2 On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

Step 3 In the **Basic Information** area of the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.

Perform this step only when Kerberos authentication is enabled for the cluster.

NOTE

- After the IAM user synchronization is complete, wait for 5 minutes before submitting a job. For details about IAM user synchronization, see [Synchronizing IAM Users to MRS..](#)
- When the policy of the user group an IAM user belongs to changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, or vice versa, it takes time for the cluster node's System Security Services Daemon (SSSD) cache to refresh. To prevent job submission failure, wait for five minutes after user synchronization is complete before submitting the job with the new policy.
- If the IAM username contains spaces (for example, **admin 01**), jobs cannot be added.

Step 4 Click **Job Management**. On the displayed job list page, click **Create**.

Step 5 In **Type**, select **MapReduce**. Configure other job information.

Figure 5-4 Adding a MapReduce job

Create Job

* Type

* Name

* Program Path


Parameters

Service Parameter

Command Reference

Table 5-1 Job configuration information

Parameter	Description	Example
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.	mapreduce_job
Program Path	<p>Path of the program package to be executed. You can enter the path or click HDFS or OBS to select a file.</p> <ul style="list-style-type: none"> The value contains a maximum of 1,023 characters. It cannot contain special characters (; &>,<'\$) and cannot be left blank or all spaces. The OBS program path should start with obs://, for example, obs://wordcount/program/XXX.jar. The HDFS program path should start with hdfs://, for example, hdfs://hacluster/user/XXX.jar. The MapReduce job execution program must end with .jar. 	obs://wordcount/program/test.jar

Parameter	Description	Example
Parameters	<p>(Optional) Key parameters for program execution. Use spaces to separate multiple parameters.</p> <p>Configuration format: <i>Program class name Data input path Data output path</i></p> <ul style="list-style-type: none"> • Program class name: It is specified by a function in your program. MRS is responsible for transferring parameters only. • Data input path: Click HDFS or OBS to select a path or manually enter a correct path. • Data output path: output path of the data processing result. Enter a directory that does not exist. The parameter contains a maximum of 150,000 characters. It cannot contain special characters ; &><'\$, but can be left blank. <p>CAUTION When entering a parameter containing sensitive information (for example, login password), you can add an at sign (@) before the parameter name to encrypt the parameter value. This prevents the sensitive information from being persisted in plaintext.</p> <p>When you view job information on the MRS console, the sensitive information is displayed as *.</p> <p>Example: <code>username=testuser @password=User password</code></p>	-
Service Parameter	<p>(Optional) Service parameters for the job.</p> <p>To modify the current job, change this parameter. For permanent changes to the entire cluster, refer to Modifying the Configuration Parameters of an MRS Cluster Component and modify the cluster component parameters accordingly.</p> <p>Click  on the right to add more parameters.</p> <p>If a job needs to access OBS using AK/SK, add the following service configuration parameters:</p> <ul style="list-style-type: none"> • fs.obs.access.key: key ID for accessing OBS. • fs.obs.secret.key: key corresponding to the key ID for accessing OBS. 	-
Command Reference	Command submitted to the background for execution when a job is submitted.	<code>yarn jar hdfs://hacluster/user/test.jar</code>

Step 6 Confirm job configuration information and click **OK**.

Step 7 After the job is submitted, you can view the job running status and execution result in the job list. After the job status changes to **Completed**, you can view the analysis result of related programs.

----End

Submitting a Job Using the Cluster Client

Step 1 Install the MRS cluster client. For details, see [Installing an MRS Cluster Client](#).

The MRS cluster comes with a client installed for job submission by default, which can also be used directly. For MRS 3.x and later versions, the default client installation path is **/opt/Bigdata/client** on the Master node. For versions earlier than MRS 3.x, the default client installation path is **/opt/client** on the Master node.

Step 2 Log in to the node where the client is located as the MRS cluster client installation user.

Step 3 Initialize environment variables.

```
cd /opt/Bigdata/client
```

```
source bigdata_env
```

Step 4 Perform authentication if Kerberos authentication has been enabled for the current cluster.

Skip this step for normal clusters.

```
kinit MRScluster service user
```

The MRS cluster service user needs to create a service user with the job submission permission on Manager. For details, see [Creating an MRS Cluster User](#).

Example:

```
kinit testuser
```

Step 5 Copy the program in the OBS file system to the node where the cluster client is located.

```
hadoop fs -Dfs.obs.access.key=AK for accessing OBS -Dfs.obs.secret.key=SK for  
accessing OBS -copyToLocal Source path of the application Destination path of  
the application
```

Example:

```
hadoop fs -Dfs.obs.access.key=XXXX -Dfs.obs.secret.key=XXXX -copyToLocal  
"obs://mrs-word/program/hadoop-mapreduce-examples-XXX.jar" "/  
home/omm/hadoop-mapreduce-examples-XXX.jar"
```

NOTE

- Commands carrying authentication passwords pose security risks. Disable historical command recording before running such commands to prevent information leakage.
- To obtain the AK and SK, log in to the OBS console and choose **My Credentials > Access Keys** from the username drop-down list in the upper right corner of the page.

- Step 6** Submit a wordcount job. If data needs to be read from OBS or outputted to OBS, the AK/SK parameters need to be added.

```
hadoop jar Application wordcount Input file path Output file path
```

Example:

```
hadoop jar /home/omm/hadoop-mapreduce-examples-XXX.jar wordcount -  
Dfs.obs.access.key=XXXX -Dfs.obs.secret.key=XXXX "obs://mrs-word/input/*"  
"obs://mrs-word/output/"
```

- *Input file path* is the path for storing job input files on OBS.
- *Output file path* is the path for storing the job output file on OBS. Set it to a directory that does not exist.

----End

5.3.2 Running a SparkSubmit Job

MRS allows you to submit and run your own programs, and get the results. This section will show you how to submit a SparkSubmit job in an MRS cluster.

Spark is an open source parallel data processing framework. It helps users easily develop unified big data applications and perform offline processing, stream processing, and interactive analysis on data.

You can create a job online and submit it for running on the MRS console, or submit a job in CLI mode on the MRS cluster client.

Prerequisites

- You have uploaded the program packages and data files required for running jobs to OBS or HDFS.
- If the job program needs to read and analyze data in the OBS file system, you need to configure storage-compute decoupling for the MRS cluster. For details, see [Configuring Storage-Compute Decoupling for an MRS Cluster](#).

Submitting a Job on the Console

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** In the **Basic Information** area of the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.

Perform this step only when Kerberos authentication is enabled for the cluster.


 **NOTE**

- After the IAM user synchronization is complete, wait for 5 minutes before submitting a job. For details about IAM user synchronization, see [Synchronizing IAM Users to MRS..](#)
- When the policy of the user group an IAM user belongs to changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, or vice versa, it takes time for the cluster node's System Security Services Daemon (SSSD) cache to refresh. To prevent job submission failure, wait for five minutes after user synchronization is complete before submitting the job with the new policy.
- If the IAM username contains spaces (for example, **admin 01**), jobs cannot be added.

Step 4 Click **Job Management**. On the displayed job list page, click **Create**.

Step 5 In **Type**, select **SparkSubmit**. Configure other job information.

Figure 5-5 Adding a Spark job



The screenshot shows a web form for adding a Spark job. The 'Type' dropdown is set to 'SparkSubmit'. The 'Name' field contains the placeholder 'Enter a job name.'. The 'Program Path' field contains 'obs://bucket/program/xx.jar or xx.py'. There are two sections for optional parameters: '(Optional) Tuning Parameters' and '(Optional) Runtime Parameters'. Each section has a 'Parameter' and 'Value' input field, a 'Create' button, and 'HDFS' and 'OBS' radio buttons. Below the runtime parameters section, there is a note: 'Use spaces to separate parameters and the '@key=value' format to prevent values being saved in plaintext.'. The '(Optional) Service Parameter' section also has 'Parameter', 'Value', and 'Create' fields. At the bottom, the 'Command Reference' is 'spark-submit --master yarn --deploy-mode cluster'.

Table 5-2 Job configuration information

Parameter	Description	Example
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.	spark_job

Parameter	Description	Example
Program Path	<p>Path of the program package to be executed. You can enter the path or click HDFS or OBS to select a file.</p> <ul style="list-style-type: none"> • The value contains a maximum of 1,023 characters. It cannot contain special characters (; &>,<'\$) and cannot be left blank or all spaces. • The OBS program path should start with obs://, for example, obs://wordcount/program/XXX.jar. The HDFS program path should start with hdfs://, for example, hdfs://hacluster/user/XXX.jar. • The SparkSubmit job execution program must end with .jar or .py. 	obs://wordcount/program/test.jar
Program Parameter	<p>(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance.</p> <p>Table 5-3 lists the common program parameters of Spark jobs. You can configure the parameters based on the execution program and cluster resources.</p>	-
Parameters	<p>(Optional) Key parameter for program execution. The parameter is specified by the function of the custom program. MRS is only responsible for loading the parameters.</p> <p>Multiple parameters are separated by spaces. The value can contain a maximum of 150,000 characters and can be left blank. The value cannot contain special characters such as ; &><'\$</p> <p>CAUTION</p> <p>When entering a parameter containing sensitive information (for example, login password), you can add an at sign (@) before the parameter name to encrypt the parameter value. This prevents the sensitive information from being persisted in plaintext.</p> <p>When you view job information on the MRS console, the sensitive information is displayed as *.</p> <p>Example: username=testuser @password=User password</p>	-

Parameter	Description	Example
Service Parameter	<p>(Optional) Service parameters for the job.</p> <p>To modify the current job, change this parameter. For permanent changes to the entire cluster, refer to Modifying the Configuration Parameters of an MRS Cluster Component and modify the cluster component parameters accordingly.</p> <p>Click to the add icon on the right to add more parameters.</p> <p>If a job needs to access OBS using AK/SK, add the following service configuration parameters:</p> <ul style="list-style-type: none"> • fs.obs.access.key: key ID for accessing OBS. • fs.obs.secret.key: key corresponding to the key ID for accessing OBS. 	-
Command Reference	Command submitted to the background for execution when a job is submitted.	spark-submit --master yarn--deploy- mode cluster

Table 5-3 Spark job running program parameters

Parameter	Description	Example
--conf	Configuration item for adding tasks.	spark.executor.memory=2G
--driver-memory	Set the running memory of driver.	2G
--num-executors	Set the number of executors to be started.	5
--executor-cores	Set the number of executor cores.	2
--class	Set the main class name of a task, which is specified by a function in the user program.	org.apache.spark.examples.SparkPi
--files	Upload files to a job. The files can be user-defined configuration files or some data files from OBS or HDFS.	-
--jars	Additional dependency JAR packages of a task, which is used to add the external dependency packages to the task.	-
--executor-memory	Executor memory.	2G

Parameter	Description	Example
<code>--conf spark-yarn.maxAppAttempts</code>	Control the number of AM retries. If this parameter is set to 0 , retry is not allowed. If this parameter is set to 1 , one retry is allowed.	0

Table 5-4 Job configuration information

Parameter	Description
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. NOTE You are advised to set different names for different jobs.
Program Path	Path of the program package to be executed. The following requirements must be met: <ul style="list-style-type: none"> The value contains a maximum of 1,023 characters. It cannot contain special characters (; &><'\$) and cannot be left blank or all spaces. The path of the program to be executed can be stored in HDFS or OBS. The path varies depending on the file system. <ul style="list-style-type: none"> OBS: The path must start with s3a://. Example: s3a://wordcount/program/xxx.jar HDFS: The path must start with /user. For details about how to import data to HDFS, see Uploading Application Data to an MRS Cluster. For SparkScript, the path must end with .sql. For MapReduce and Spark, the path must end with .jar. The .sql and .jar are case-insensitive.
Parameters	Key parameter for program execution. The parameter is specified by the function of the user's program. MRS is only responsible for loading the parameter. Multiple parameters are separated by space. Configuration method: <i>Package name.Class name</i> The parameter contains a maximum of 150,000 characters. It cannot contain special characters ; &><'\$, but can be left blank. CAUTION When entering a parameter containing sensitive information (for example, login password), you can add an at sign (@) before the parameter name to encrypt the parameter value. This prevents the sensitive information from being persisted in plaintext. When you view job information on the MRS console, the sensitive information is displayed as *. Example: username=testuser @password=User password

Parameter	Description
Import From	<p>Path for inputting data</p> <p>Data can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> • OBS: The path must start with s3a://. • HDFS: The path must start with /user. For details about how to import data to HDFS, see Uploading Application Data to an MRS Cluster. <p>The value contains a maximum of 1,023 characters and cannot contain special characters (; &>,<'\$). This parameter can be left blank.</p>
Export To	<p>Path for outputting data</p> <p>NOTE</p> <ul style="list-style-type: none"> • When setting this parameter, select OBS or HDFS. Select a file directory or manually enter a file directory, and click OK. • If you add the hadoop-mapreduce-examples-x.x.x.jar sample program or a program similar to hadoop-mapreduce-examples-x.x.x.jar, enter a directory that does not exist. <p>Data can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> • OBS: The path must start with s3a://. • HDFS: The path must start with /user. <p>The value contains a maximum of 1,023 characters and cannot contain special characters (; &>,<'\$). This parameter can be left blank.</p>
Log Path	<p>Path for storing job logs that record job running status.</p> <p>Data can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> • OBS: The path must start with s3a://. • HDFS: The path must start with /user. <p>The value contains a maximum of 1,023 characters and cannot contain special characters (; &>,<'\$). This parameter can be left blank.</p>

Step 6 Confirm job configuration information and click **OK**.

Step 7 After the job is submitted, you can view the job running status and execution result in the job list. After the job status changes to **Completed**, you can view the analysis result of related programs.

----End

Submitting a Job Using the Cluster Client

Step 1 Install the MRS cluster client. For details, see [Installing an MRS Cluster Client](#).

The MRS cluster comes with a client installed for job submission by default, which can also be used directly. For MRS 3.x and later versions, the default client installation path is `/opt/Bigdata/client` on the Master node. For versions earlier than MRS 3.x, the default client installation path is `/opt/client` on the Master node.

Step 2 Upload the application to be run to the node where the cluster client is located.

 **NOTE**

- The JAR sample program used in this section is `{Cluster client installation directory}/Spark2x/spark/examples/jars/spark-examples_*.jar`. (In some versions, the name of the `Spark2x` folder in the cluster is `Spark`. Replace it with the actual name.)
- You can log in to the client node and run the following command to upload the sample JAR package to be executed to the HDFS.

For example, upload the file to the `/tmp` directory of HDFS.

```
hdfs dfs -put {Client installation directory}/Spark2x/spark/examples/jars/spark-examples_*.jar /tmp
```

Step 3 If Kerberos authentication has been enabled for the current cluster, create a user for submitting jobs by referring to [Creating an MRS Cluster User](#).

Skip this step for normal clusters.

In this example, a machine-machine user has been created, and user groups (`hadoop` and `supergroup`), the primary group (`supergroup`), and role permissions (`System_administrator` and `default`) have been correctly assigned to the user.

After the user is created, download the authentication credential file.

- For clusters of MRS 3.x or later, log in to FusionInsight Manager and choose **System > Permission > User**. In the **Operation** column of the newly created user, choose **More > Download Authentication Credential**.
- For MRS 2.x or earlier, log in to MRS Manager and choose **System > Manage User**. In the **Operation** column of the newly created user, choose **More > Download Authentication Credential**.

Upload the user authentication credential to the `/opt` directory on the cluster client node.

Step 4 Log in to the node where the client is located as the MRS cluster client installation user.

Step 5 Decompress the user authentication credential file to obtain the `user.keytab` and `krb5.conf` files.

```
cd /opt
tar -xvf XXX_keytab.tar
```

Step 6 Initialize environment variables.

```
cd /opt/Bigdata/client
source bigdata_env
cd $SPARK_HOME
```

Step 7 Run the following command to submit a Spark job:

```
./bin/spark-submit --master yarn --deploy-mode client --conf  
spark.yarn.principal=MRSTest --conf spark.yarn.keytab=/opt/user.keytab --class  
org.apache.spark.examples.SparkPi examples/jars/spark-examples_*.jar 10
```

```
...  
Pi is roughly 3.1402231402231404
```

Parameter description:

- **deploy-mode**: running mode of the Spark driver. The value can be **client** or **cluster**. In this example, the client mode is used to submit the job.
- **conf**: additional configuration attribute. For example, if the **keytab** file is used for user authentication in this example, configure the following parameters:
 - **spark.yarn.principal**: name of the user who submits the job.
 - **spark.yarn.keytab**: **keytab** file for user authentication.
- **class**: main class name of the application, which is specified by the running application.
- **XXX.jar**: program run by the job

----End

5.3.3 Running a HiveSQL Job

MRS allows you to submit and run your own programs, and get the results. This section will show you how to submit a HiveSQL job in an MRS cluster.

HiveSQL jobs enable the submission of SQL statements and script files for data analysis and querying. Both SQL statements and scripts are supported, and script files can be used to submit sensitive information.

You can create a job online and submit it for running on the MRS console, or submit a job in CLI mode on the MRS cluster client.

Prerequisites

- You have uploaded the program packages and data files required by jobs to OBS or HDFS.
- If the job program needs to read and analyze data in the OBS file system, you need to configure storage-compute decoupling for the MRS cluster. For details, see [Configuring Storage-Compute Decoupling for an MRS Cluster](#).

Submitting a Job on the Console

Step 1 Log in to the MRS console.

Step 2 On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

Step 3 In the **Basic Information** area of the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.

Perform this step only when Kerberos authentication is enabled for the cluster.

NOTE

- After the IAM user synchronization is complete, wait for 5 minutes before submitting a job. For details about IAM user synchronization, see [Synchronizing IAM Users to MRS..](#)
- When the policy of the user group an IAM user belongs to changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, or vice versa, it takes time for the cluster node's System Security Services Daemon (SSSD) cache to refresh. To prevent job submission failure, wait for five minutes after user synchronization is complete before submitting the job with the new policy.
- If the IAM username contains spaces (for example, **admin 01**), jobs cannot be added.

Step 4 Click **Job Management**. On the displayed job list page, click **Create**.

Step 5 Set **Type** to **HiveSql** and configure HiveSQL job information by referring to [Table 5-5](#).

Figure 5-6 Adding a HiveSQL job

Table 5-5 Job configuration information

Parameter	Description	Example
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.	hivesql

Parameter	Description	Example
SQL Type	<p>Submission type of the SQL statement.</p> <ul style="list-style-type: none"> SQL: Run the entered SQL statement. Script: Load SQL scripts from HDFS or OBS to run SQL statements. 	SQL
SQL Statement	<p>This parameter is valid only when SQL Type is set to SQL. Enter the SQL statement to be executed, and then click Check to check whether the SQL statement is correct.</p> <p>If you want to submit and execute multiple statements at the same time, use semicolons (;) to separate them.</p>	-
SQL File	<p>This parameter is valid only when SQL Type is set to Script. The path of the SQL file to be executed must meet the following requirements:</p> <p>Path of the SQL script file to be executed. You can enter the path or click HDFS or OBS to select a file.</p> <ul style="list-style-type: none"> The value contains a maximum of 1,023 characters. It cannot contain special characters (; &><'\$) and cannot be left blank or all spaces. The OBS program path should start with obs://, for example, obs://wordcount/program/XXX.jar. The HDFS program path should start with hdfs://, for example, hdfs://hacluster/user/XXX.jar. The HiveScript file must end with .sql. 	obs://wordcount/program/test.sql
Program Parameter	<p>(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance.</p> <p>Table 5-6 lists the common program parameters of HiveSQL jobs. You can configure the parameters based on the execution program and cluster resources.</p>	-


Parameter	Description	Example
Service Parameter	<p>(Optional) Service parameters for the job.</p> <p>To modify the current job, change this parameter. For permanent changes to the entire cluster, refer to Modifying the Configuration Parameters of an MRS Cluster Component and modify the cluster component parameters accordingly.</p> <p>Click  on the right to add more parameters.</p> <p>For example, add the following service configuration parameters:</p> <ul style="list-style-type: none"> • fs.obs.access.key: key ID used to access OBS in AK/SK mode. • fs.obs.secret.key: key used to access OBS in AK/SK mode. • hive.execution.engine: engine for executing the job. The value can be mr or tez. 	-
Command Reference	Command submitted to the background for execution when a job is submitted.	beeline- e"SELECT * for TABLE test;"

Table 5-6 HiveSQL job running program parameters

Parameter	Description	Example Value
--hiveconf	Configure the Hive service.	<p>For example, set the execution engine to MapReduce.</p> <ul style="list-style-type: none"> • Parameter: --hiveconf • Value: hive.execution.engine=mr
--hivevar	Set user-defined variables.	<p>Set the variable ID as follows:</p> <ul style="list-style-type: none"> • Parameter: --hivevar id • Value: "123" select * from test where id = \$ {hivevar:id};

Step 6 Confirm job configuration information and click **OK**.

Step 7 After the job is submitted, you can view the job running status and execution result in the job list. After the job status changes to **Completed**, you can view the analysis result of related programs.

----End

Submitting a Job Using the Cluster Client

Step 1 Install the MRS cluster client. For details, see [Installing an MRS Cluster Client](#).

The MRS cluster comes with a client installed for job submission by default, which can also be used directly. For MRS 3.x and later versions, the default client installation path is `/opt/Bigdata/client` on the Master node. For versions earlier than MRS 3.x, the default client installation path is `/opt/client` on the Master node.

Step 2 Log in to the node where the client is located as the MRS cluster client installation user.

Step 3 Initialize environment variables.

```
cd /opt/Bigdata/client
```

```
source bigdata_env
```

Step 4 Perform authentication if Kerberos authentication has been enabled for the current cluster.

Skip this step for normal clusters.

```
kinit MRScluster service user
```

The MRS cluster service user needs to create a service user who has the permission to submit jobs on Manager. The user needs to be added to the hive user group. For details, see [Creating an MRS Cluster User](#).

Example:

```
kinit testuser
```

Step 5 Run the **beeline** command to connect to Hive in the cluster and run related tasks.

```
beeline -f SQL file (SQL statements in the execution files.)
```

To specify a component service user for normal clusters, use the command below. If no service user is specified, HiveServer will connect using the current OS user.

```
beeline -nMRS cluster service user
```

----End

5.3.4 Running a Spark SQL Job

MRS allows you to submit and run your own programs, and get the results. This section will show you how to submit a Spark SQL job in an MRS cluster.

Spark SQL jobs are used to query and analyze data, including SQL statements and scripts. If SQL statements contain sensitive information, you can also use script files to submit them.

You can create a job online and submit it for running on the MRS console, or submit a job in CLI mode on the MRS cluster client.

Prerequisites

- You have uploaded the program packages and data files required by jobs to OBS or HDFS.
- If the job program needs to read and analyze data in the OBS file system, you need to configure storage-compute decoupling for the MRS cluster. For details, see [Configuring Storage-Compute Decoupling for an MRS Cluster](#).

Submitting a Job on the Console

Step 1 Log in to the MRS console.

Step 2 On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

Step 3 In the **Basic Information** area of the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.

Perform this step only when Kerberos authentication is enabled for the cluster.


NOTE

- After the IAM user synchronization is complete, wait for 5 minutes before submitting a job. For details about IAM user synchronization, see [Synchronizing IAM Users to MRS](#).
- When the policy of the user group an IAM user belongs to changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, or vice versa, it takes time for the cluster node's System Security Services Daemon (SSSD) cache to refresh. To prevent job submission failure, wait for five minutes after user synchronization is complete before submitting the job with the new policy.
- If the IAM username contains spaces (for example, **admin 01**), jobs cannot be added.

Step 4 Click **Job Management**. On the displayed job list page, click **Create**.

Step 5 Set **Type** to **SparkSql** and configure Spark SQL information by referring to [Table 5-7](#).

Figure 5-7 Adding a Spark SQL job



The screenshot shows the configuration page for adding a Spark SQL job. It includes the following elements:

- Type:** A dropdown menu set to "SparkSql".
- Name:** A text input field with the placeholder "Enter a job name."
- SQL Type:** Two buttons, "SQL" (highlighted in blue) and "Script".
- SQL Statement:** A large text area for entering the SQL statement, with a "Check" button to the right.
- (Optional) Tuning Parameters:** A table with columns "Parameter" and "Value", and a "Create" button.
- (Optional) Service Parameter:** A table with columns "Parameter" and "Value", and a "Create" button.
- Command Reference:** A text box containing the command `mrs-spark-sql-wrapper -e`.

Table 5-7 Job configuration information

Parameter	Description	Example
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.	sparksql
SQL Type	Submission type of the SQL statement <ul style="list-style-type: none"> SQL: Run the entered SQL statement. Script: Load SQL scripts from HDFS or OBS to run SQL statements. 	SQL
SQL Statement	This parameter is valid only when SQL Type is set to SQL . Enter the SQL statement to be executed, and then click Check to check whether the SQL statement is correct. If you want to submit and execute multiple statements at the same time, use semicolons (;) to separate them.	-

Parameter	Description	Example
SQL File	<p>This parameter is valid only when SQL Type is set to Script. The path of the SQL file to be executed must meet the following requirements:</p> <p>Path of the SQL script file to be executed. You can enter the path or click HDFS or OBS to select a file.</p> <ul style="list-style-type: none"> • The value contains a maximum of 1,023 characters. It cannot contain special characters (; &>,<'\$) and cannot be left blank or all spaces. • The OBS program path should start with obs://, for example, obs://wordcount/program/XXX.jar. The HDFS program path should start with hdfs://, for example, hdfs://hacluster/user/XXX.jar. • The script file must end with .sql. 	obs://wordcount/program/test.sql
Program Parameter	<p>(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance.</p> <p>Table 5-8 lists the common program parameters of SparkSql jobs. You can configure the parameters based on the execution program and cluster resources.</p>	-
Service Parameter	<p>(Optional) Service parameters for the job.</p> <p>To modify the current job, change this parameter. For permanent changes to the entire cluster, refer to Modifying the Configuration Parameters of an MRS Cluster Component and modify the cluster component parameters accordingly.</p> <p>Click to the add icon on the right to add more parameters.</p> <p>For example, add the following service configuration parameters:</p> <ul style="list-style-type: none"> • fs.obs.access.key: key ID used to access OBS in AK/SK mode. • fs.obs.secret.key: key used to access OBS in AK/SK mode. 	-
Command Reference	Command submitted to the background for execution when a job is submitted.	mrs-spark-sql-wrapper -e

Table 5-8 Program parameters

Parameter	Description	Example
--conf	Configuration item for adding tasks.	spark.executor.memory=2G
--driver-memory	Running memory of a driver.	2G
--num-executors	Number of executors to be started.	5
--executor-cores	Number of executor cores.	2
--jars	Additional dependency packages of a task, which is used to add the external dependency packages to the task.	-
--executor-memory	Executor memory.	2G

Step 6 Confirm job configuration information and click **OK**.

Step 7 After the job is submitted, you can view the job running status and execution result in the job list. After the job status changes to **Completed**, you can view the analysis result of related programs.

----End

Submitting a Job Using the Cluster Client

Step 1 Install the MRS cluster client. For details, see [Installing an MRS Cluster Client](#).

The MRS cluster comes with a client installed for job submission by default, which can also be used directly. For MRS 3.x and later versions, the default client installation path is `/opt/Bigdata/client` on the Master node. For versions earlier than MRS 3.x, the default client installation path is `/opt/client` on the Master node.

Step 2 If Kerberos authentication has been enabled for the current cluster, create a user for submitting jobs by referring to [Creating an MRS Cluster User](#).

Skip this step for normal clusters.

In this example, a machine-machine user has been created, and user groups (**hadoop** and **supergroup**), the primary group (**supergroup**), and role permissions (**System_administrator** and **default**) have been correctly assigned to the user.

After the user is created, download the authentication credential file.

- For clusters of MRS 3.x or later, log in to FusionInsight Manager and choose **System > Permission > User**. In the **Operation** column of the newly created user, choose **More > Download Authentication Credential**.
- For MRS 2.x or earlier, log in to MRS Manager and choose **System > Manage User**. In the **Operation** column of the newly created user, choose **More > Download Authentication Credential**.

Upload the user authentication credential to the `/opt` directory on the cluster client node.

Step 3 Log in to the node where the client is located as the MRS cluster client installation user.

Step 4 Decompress the user authentication credential file to obtain the **user.keytab** and **krb5.conf** files.

```
cd /opt
tar -xvf XXX_keytab.tar
```

Step 5 Initialize environment variables.

```
cd /opt/Bigdata/client
source bigdata_env
cd $SPARK_HOME
```

Step 6 Enter the spark-sql CLI and run the SQL statement.

```
./bin/spark-sql --conf spark.yarn.principal=MRSTest --conf
spark.yarn.keytab=/opt/user.keytab
```

To execute the SQL file, upload the SQL file to the node where the client is located (for example, the **/opt/** directory) in advance and run the following command:

```
./bin/spark-sql --conf spark.yarn.principal=MRSTest --conf
spark.yarn.keytab=/opt/user.keytab -f /opt/script.sql
```

- **spark.yarn.principal**: name of the user who submits the job.
- **spark.yarn.keytab**: **keytab** file for user authentication.

----End

5.3.5 Running a Flink Job

MRS allows you to submit and run your own programs, and get the results. This section will show you how to submit a Flink job in an MRS cluster.

Flink jobs are used to submit JAR programs to process streaming data.

You can create a job online and submit it for running on the MRS console, or submit a job in CLI mode on the MRS cluster client.

Prerequisites

- You have uploaded the program packages and data files required by jobs to OBS or HDFS.
- If the job program needs to read and analyze data in the OBS file system, you need to configure storage-compute decoupling for the MRS cluster. For details, see [Configuring Storage-Compute Decoupling for an MRS Cluster](#).

Submitting a Job on the Console

Step 1 Log in to the MRS console.

Step 2 On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

Step 3 In the **Basic Information** area of the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.

Perform this step only when Kerberos authentication is enabled for the cluster.

NOTE

- After the IAM user synchronization is complete, wait for 5 minutes before submitting a job. For details about IAM user synchronization, see [Synchronizing IAM Users to MRS..](#)
- When the policy of the user group an IAM user belongs to changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, or vice versa, it takes time for the cluster node's System Security Services Daemon (SSSD) cache to refresh. To prevent job submission failure, wait for five minutes after user synchronization is complete before submitting the job with the new policy.
- If the IAM username contains spaces (for example, **admin 01**), jobs cannot be added.

Step 4 Click **Job Management**. On the displayed job list page, click **Create**.

Step 5 Set **Type** to **Flink** and configure Flink job information by referring to [Table 5-9](#).

Figure 5-8 Creating a Flink job

The screenshot shows a 'Create Job' form with the following fields and options:

- Type:** A dropdown menu with 'Flink' selected.
- Name:** A text input field with the placeholder 'Enter a job name.'
- Program Path:** A text input field with the placeholder 'obs://bucket/program/xx.jar'. To its right are two buttons: 'HDFS' and 'OBS'.
- (Optional) Tuning Parameters:** A section with a help icon. It contains two input fields: 'Parameter' and 'Value', followed by a 'Create' button.
- (Optional) Runtime Parameters:** A section with a help icon. It contains a large empty text area for parameters. To its right are two buttons: 'HDFS' and 'OBS'.
- (Optional) Service Parameter:** A section with a help icon. It contains two input fields: 'Parameter' and 'Value', followed by a 'Create' button.
- Command Reference:** A text field containing the command 'flink run -d -m yarn-cluster'.

Table 5-9 Job configuration information

Parameter	Description	Example
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.	flink_job

Parameter	Description	Example
Program Path	<p>Path of the program package to be executed. You can enter the path or click HDFS or OBS to select a file.</p> <ul style="list-style-type: none"> The value contains a maximum of 1,023 characters. It cannot contain special characters (; &>,<'\$) and cannot be left blank or all spaces. The OBS program path should start with obs://, for example, obs://wordcount/program/XXX.jar. The HDFS program path should start with hdfs://, for example, hdfs://hacluster/user/XXX.jar. The Flink job execution program must end with .jar. 	-
Program Parameter	<p>(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance.</p> <p>Table 5-10 lists the common program parameters of Flink jobs. You can configure the parameters based on the execution program and cluster resources.</p>	-
Parameters	<p>(Optional) Key parameter for program execution. The parameter is specified by the function of the custom program. MRS is only responsible for loading the parameters.</p> <p>Multiple parameters are separated by spaces. The value can contain a maximum of 150,000 characters and can be left blank. The value cannot contain special characters such as ; &><'\$</p> <p>CAUTION</p> <p>When entering a parameter containing sensitive information (for example, login password), you can add an at sign (@) before the parameter name to encrypt the parameter value. This prevents the sensitive information from being persisted in plaintext.</p> <p>When you view job information on the MRS console, the sensitive information is displayed as *.</p> <p>Example: username=testuser @password=User password</p>	-


Parameter	Description	Example
Service Parameter	<p>(Optional) Service parameters for the job.</p> <p>To modify the current job, change this parameter. For permanent changes to the entire cluster, refer to Modifying the Configuration Parameters of an MRS Cluster Component and modify the cluster component parameters accordingly.</p> <p>Click  on the right to add more parameters.</p> <p>If a job needs to access OBS using AK/SK, add the following service configuration parameters:</p> <ul style="list-style-type: none"> • fs.obs.access.key: key ID for accessing OBS. • fs.obs.secret.key: key corresponding to the key ID for accessing OBS. 	-
Command Reference	Command submitted to the background for execution when a job is submitted.	-

Table 5-10 Flink job running program parameters

Parameter	Description	Example
-ytm	Memory size of each TaskManager container. (Optional unit. The unit is MB by default.)	1024
-yjm	Memory size of JobManager container. (Optional unit. The unit is MB by default.)	1024
-ys	Number of TaskManager cores.	2
-ynm	Custom name of an application on YARN.	test
-c	Class of the program entry method (for example, the main or getPlan() method). This parameter is required only when the JAR program does not specify the class of its manifest.	com.bigdata.mrs.test

Step 6 Confirm job configuration information and click **OK**.

Step 7 After the job is submitted, you can view the job running status and execution result in the job list. After the job status changes to **Completed**, you can view the analysis result of related programs.

----End

Submitting a Job Using the Cluster Client

Step 1 Install the MRS cluster client. For details, see [Installing an MRS Cluster Client](#).

The MRS cluster comes with a client installed for job submission by default, which can also be used directly. For MRS 3.x and later versions, the default client installation path is `/opt/Bigdata/client` on the Master node. For versions earlier than MRS 3.x, the default client installation path is `/opt/client` on the Master node.

Step 2 Log in to the node where the client is located as the MRS cluster client installation user.

Step 3 Run the following command to initialize environment variables:

```
cd /opt/Bigdata/client
source bigdata_env
```

Step 4 Perform the following steps to create a user for submitting jobs and modify the security configuration of the cluster client only for clusters with Kerberos authentication enabled.

1. Prepare a user for submitting Flink jobs.
 - MRS 3.x or earlier: For details, see [Preparing a Development User](#).
 - MRS 3.x or later: For details, see [Preparing a Development User](#).
2. Log in to Manager as the newly created user.
 - For MRS 3.x earlier: Log in to Manager of the cluster. Choose **System > Manage User**. In the **Operation** column of the row that contains the added user, choose **More > Download authentication credential** to locate the row that contains the user.
 - For MRS 3.x or later: Log in to Manager of the cluster. Choose **System > Permission > Manage User**. On the displayed page, locate the row that contains the added user, click **More** in the **Operation** column, and select **Download authentication credential**.
3. Decompress the downloaded authentication credential package and copy the obtained file to a directory on the client node, for example, `/opt/Bigdata/client/Flink/flink/conf`. If the client is installed on a node outside the cluster, copy the obtained files to the `/etc/` directory on the node.
4. For MRS 3.x or later: In security mode, add the service IP address of the node where the client is installed and floating IP address of Manager to the `jobmanager.web.allow-access-address` configuration item in the `/opt/Bigdata/client/Flink/flink/conf/flink-conf.yaml` file.
5. Run the following commands to configure security authentication by adding the `keytab` path and username to the `/opt/Bigdata/client/Flink/flink/conf/flink-conf.yaml` configuration file.

```
security.kerberos.login.keytab: <user.keytab file path>
security.kerberos.login.principal: <Username>
```

Example:

```
security.kerberos.login.keytab: /opt/Bigdata/client/Flink/flink/conf/user.keytab
security.kerberos.login.principal: test
```
6. In the `bin` directory of the Flink client, run the following command to perform security hardening. Then, set a password for submitting jobs.

sh generate_keystore.sh

This script automatically replaces the SSL value in the `/opt/Bigdata/client/Flink/flink/conf/flink-conf.yaml` file. For MRS 3.x or earlier, external SSL is disabled by default in security clusters. To enable external SSL, run this script again after configuration. The configuration parameters do not exist in the default Flink configuration of MRS, if you enable SSL for external connections, you need to add the parameters listed in [Table 5-11](#).

Table 5-11 Parameter description

Parameter	Description	Example Value
security.ssl.rest.enabled	Switch to enable external SSL.	true
security.ssl.rest.keystore	Path for storing keystore .	<code>\${path}/flink.keystore</code>
security.ssl.rest.keystore-password	Password of the keystore . 123456 indicates a user-defined password is required.	123456
security.ssl.rest.key-password	Password of the SSL key. 123456 indicates a user-defined password is required.	123456
security.ssl.rest.truststore	Path for storing the truststore .	<code>\${path}/flink.truststore</code>
security.ssl.rest.truststore-password	Password of the truststore . 123456 indicates a user-defined password is required.	123456

NOTE

- For MRS 3.x or earlier: The **generate_keystore.sh** script is automatically generated.
- The generated **flink.keystore**, **flink.truststore**, and **security.cookie** items are automatically filled in the corresponding configuration items in **flink-conf.yaml**.
- For MRS 3.x or later: You can obtain the values of **security.ssl.key-password**, **security.ssl.keystore-password**, and **security.ssl.truststore-password** using the Manager plaintext encryption API by running the following command:

```
curl -k -i -u <user name>:<password> -X POST -HContent-type:application/json -d '{"plainText":"<password>"}' 'https://x.x.x.x:28443/web/api/v2/tools/encrypt';
```

In the preceding command, `<password>` must be the same as the password used for issuing the certificate, and `x.x.x.x` indicates the floating IP address of Manager in the cluster.

Commands carrying authentication passwords pose security risks. Disable historical command recording before running such commands to prevent information leakage.

7. Configure paths for the client to access the **flink.keystore** and **flink.truststore** files.
 - Absolute path: After the script is executed, the file path of **flink.keystore** and **flink.truststore** is automatically set to the absolute path **opt/**

- Bigdata/client/Flink/flink/conf/** in the **flink-conf.yaml** file. In this case, you need to move the **flink.keystore** and **flink.truststore** files from the **conf** directory to this absolute path on the Flink client and YARN nodes.
- Relative path: Perform the following steps to set the file path of **flink.keystore** and **flink.truststore** to the relative path and ensure that the directory where the Flink client command is executed can directly access the relative paths.
 - i. In the **/opt/Bigdata/client/Flink/flink/conf/** directory, create a new directory, for example, **ssl**.
 - ii. Move the **flink.keystore** and **flink.truststore** file to the **/opt/Bigdata/client/Flink/flink/conf/ssl/** directory.
 - iii. For MRS 3.x or later: Change the values of the following parameters in the **flink-conf.yaml** file to relative paths:

```
security.ssl.keystore: ssl/flink.keystore
security.ssl.truststore: ssl/flink.truststore
```
 - iv. For MRS 3.x or earlier: Change the values of the following parameters in the **flink-conf.yaml** file to relative paths:

```
security.ssl.internal.keystore: ssl/flink.keystore
security.ssl.internal.truststore: ssl/flink.truststore
```
8. If the client is installed on a node outside the cluster, add the following configuration to the configuration file (for example, **/opt/Bigdata/client/Flink/flink/conf/flink-conf.yaml**). Replace **xx.xx.xxx.xxx** with the IP address of the node where the client resides.
- ```
web.access-control-allow-origin: xx.xx.xxx.xxx
jobmanager.web.allow-access-address: xx.xx.xxx.xxx
```

### Step 5 Run the Flink job.

This section uses the WordCount sample program provided by the client as an example.

- Normal cluster (Kerberos authentication disabled)
  - Run the following commands to start a session and submit a job in the session:

```
yarn-session.sh -nm "session-name" -d
flink run /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar
```
  - Run the following command to submit a single job on YARN:

```
flink run -m yarn-cluster /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar
```
- Security cluster (Kerberos authentication enabled)
  - If the **flink.keystore** and **flink.truststore** file are stored in the absolute path:
    - Run the following commands to start a session and submit a job in the session:

```
yarn-session.sh -nm "session-name" -d
flink run /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar
```
    - Run the following command to submit a single job on YARN:

```
flink run -m yarn-cluster /opt/Bigdata/client/Flink/flink/
examples/streaming/WordCount.jar
```

- If the **flink.keystore** and **flink.truststore** file are stored in the relative path:
  - In the same directory of SSL, run the following command to start a session and submit jobs in the session. The SSL directory is a relative path. For example, if the SSL directory is **opt/Bigdata/client/Flink/flink/conf/**, then run the following command in this directory:

```
yarn-session.sh -t ssl/ -nm "session-name" -d
flink run /opt/Bigdata/client/Flink/flink/examples/streaming/
WordCount.jar
```
  - Run the following command to submit a single job on YARN:

```
flink run -m yarn-cluster -yt ssl/ /opt/Bigdata/client/Flink/flink/
examples/streaming/WordCount.jar
```

----End

### 5.3.6 Running a HadoopStreaming Job

MRS allows you to submit and run your own programs, and get the results. This section will show you how to submit a Hadoop Streaming job in an MRS cluster.

#### Prerequisites

- You have uploaded the program packages and data files required by jobs to OBS or HDFS.
- If the job program needs to read and analyze data in the OBS file system, you need to configure storage-compute decoupling for the MRS cluster. For details, see [Configuring Storage-Compute Decoupling for an MRS Cluster](#).

#### Submitting a Hadoop Streaming job

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** In the **Basic Information** area of the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.

Perform this step only when Kerberos authentication is enabled for the cluster.


#### NOTE

- After the IAM user synchronization is complete, wait for 5 minutes before submitting a job. For details about IAM user synchronization, see [Synchronizing IAM Users to MRS..](#)
- When the policy of the user group an IAM user belongs to changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, or vice versa, it takes time for the cluster node's System Security Services Daemon (SSSD) cache to refresh. To prevent job submission failure, wait for five minutes after user synchronization is complete before submitting the job with the new policy.
- If the IAM username contains spaces (for example, **admin 01**), jobs cannot be added.

**Step 4** Click **Job Management**. On the displayed job list page, click **Create**.

**Step 5** Set **Type** to **HadoopStreaming**. Configure job information by referring to [Table 5-12](#).

**Table 5-12** Job parameters

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Example    |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| Name              | Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | hadoop_job |
| Program Parameter | (Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance.<br><a href="#">Table 5-13</a> describes the common parameters of a running program.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | -          |
| Parameters        | (Optional) Key parameter for program execution. The parameter is specified by the function of the custom program. MRS is only responsible for loading the parameters.<br><br>Multiple parameters are separated by spaces. The value can contain a maximum of 150,000 characters and can be left blank. The value cannot contain special characters such as ; &><'\$<br><b>CAUTION</b><br>When entering a parameter containing sensitive information (for example, login password), you can add an at sign (@) before the parameter name to encrypt the parameter value. This prevents the sensitive information from being persisted in plaintext.<br><br>When you view job information on the MRS console, the sensitive information is displayed as *.<br><br>Example: <code>username=testuser @password=User password</code> | -          |
| Service Parameter | (Optional) Service parameters for the job.<br>To modify the current job, change this parameter. For permanent changes to the entire cluster, refer to <a href="#">Modifying the Configuration Parameters of an MRS Cluster Component</a> and modify the cluster component parameters accordingly.<br><br>Click  on the right to add more parameters.<br>If a job needs to access OBS using AK/SK, add the following service configuration parameters: <ul style="list-style-type: none"> <li><b>fs.obs.access.key</b>: key ID for accessing OBS.</li> <li><b>fs.obs.secret.key</b>: key corresponding to the key ID for accessing OBS.</li> </ul>                                                                                            | -          |



| Parameter         | Description                                                     | Example |
|-------------------|-----------------------------------------------------------------|---------|
| Command Reference | Commands submitted to the background when the job is submitted. | -       |

**Table 5-13** Program parameters

| Parameter | Description                                                                                                                                                                                | Example Value        |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| -ytm      | Memory size of each TaskManager container. (Optional unit. The unit is MB by default.)                                                                                                     | 1024                 |
| -yjm      | Memory size of JobManager container. (Optional unit. The unit is MB by default.)                                                                                                           | 1024                 |
| -yn       | Number of Yarn containers allocated to applications. The value is the same as the number of TaskManagers.<br>For MRS 3.x or later, the <b>-yn</b> parameter is not supported.              | 2                    |
| -ys       | Number of TaskManager cores                                                                                                                                                                | 2                    |
| -ynm      | Custom name of an application on Yarn                                                                                                                                                      | test                 |
| -c        | Class of the program entry method (for example, the <b>main</b> or <b>getPlan()</b> method). This parameter is required only when the JAR file does not specify the class of its manifest. | com.bigdata.mrs.test |

**Step 6** Confirm job configuration information and click **OK**.

After the job is created, you can manage it.

----End

## 5.4 Viewing MRS Job Details and Logs

You can check the status, configuration, and run logs of all jobs in an MRS cluster on the console.

Spark SQL and DistCp jobs do not generate background logs, which means that you cannot access the running logs of these jobs online.

## Viewing the Job Status

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** Click **Job Management** to view the list and status of jobs created in an MRS cluster.

By default, jobs in the job list are sorted by submission time. For details about the parameters in the job list, see [Table 5-14](#). You can quickly filter jobs by job type and job status.

**Table 5-14** Job list parameters

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name/ID    | Job name, which is set when a job is created.<br>ID is the unique identifier of a job. After a job is added, the system automatically assigns a value to ID.                                                                                                                                                                                                                                                                                                                                                           |
| Username   | Name of the user who submits a job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Job Type   | Job type.<br><b>NOTE</b><br>After importing and exporting files on the <b>Files</b> tab page, you can view the DistCp job on the <b>Jobs</b> page.                                                                                                                                                                                                                                                                                                                                                                     |
| Status     | Job status. <ul style="list-style-type: none"><li>• <b>Submitted</b>: The job has been submitted.</li><li>• <b>Accepted</b>: Initial status of a job after it is successfully submitted.</li><li>• <b>Running</b>: The job is running.</li><li>• <b>Completed</b>: The job is executed.</li><li>• <b>Terminated</b>: The job execution is stopped.</li><li>• <b>Abnormal</b>: An error is reported during job execution, or the job execution is complete but fails.</li></ul>                                         |
| Result     | Execution result of a job. <ul style="list-style-type: none"><li>• <b>Undefined</b>: indicates that the job is being executed.</li><li>• <b>Successful</b>: indicates that the job has been executed.</li><li>• <b>Killed</b>: indicates that the job is manually terminated during execution.</li><li>• <b>Failed</b>: indicates that the job fails to be executed.</li></ul> <b>NOTE</b><br>After a job is successfully executed or fails to be executed, it cannot be executed again. You can only add a job again. |
| Queue Name | Name of the resource queue bound to the user who submits the job.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Submitted  | Time when a job is submitted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

| Parameter | Description                                                                                          |
|-----------|------------------------------------------------------------------------------------------------------|
| Ended     | Time when a job is completed or manually stopped.                                                    |
| Operation | Operations can be performed on the job. For details, see <a href="#">Managing MRS Cluster Jobs</a> . |

----End

## Viewing MRS Job Logs

**Step 1** Log in to the MRS management console.

**Step 2** Choose **Active Clusters** in the navigation pane, select a running cluster, and click its name to switch to the cluster details page.

**Step 3** Click **Jobs**.

**Step 4** In the **Operation** column of the job to be viewed, choose **More > View Details**.

In the **View Details** window that is displayed, configuration of the selected job is displayed.

**Step 5** Select a running job, and click **View Log** in the **Operation** column.

In the new page that is displayed, real-time log information of the job is displayed.

### NOTE

- Each tenant can submit and view 10 jobs concurrently.
- After a job is executed, the system will compress and save the logs to the corresponding path if you choose to save job logs to OBS or HDFS. Therefore, after a job execution of this type is completed, the job status is still **Running**. After the log is successfully stored, the job status changes to **Completed**. The log storage duration depends on the log size and takes several minutes.

----End

# 6 Managing Clusters

## 6.1 Overview

### MRS Console

After creating a cluster, you can use the MRS console or MRS Manager to manage and maintain it, as well as view basic information.

- The MRS console's cluster management page allows you to monitor and manage nodes, components, alarms, files, and jobs related to the cluster.
- MRS Manager is the O&M management system of MRS and provides unified cluster management capabilities for services deployed in clusters.

For further information on the capabilities of the MRS console and Manager, refer to [Table 6-1](#).

**Table 6-1** MRS console and Manager capability comparison

| Operation                                                                                                                  | Console | Manager |
|----------------------------------------------------------------------------------------------------------------------------|---------|---------|
| Changing subnets, adding security group rules, controlling OBS permissions, managing agencies, and synchronizing IAM users | Yes     | No      |
| Adding node groups, scaling out, scaling in, and upgrading specifications                                                  | Yes     | No      |
| Isolating hosts, starting all roles, and stopping all roles                                                                | Yes     | Yes     |
| Downloading the client, starting services, stopping services, and perform rolling restart of services                      | Yes     | Yes     |

| Operation                                                                                         | Console | Manager |
|---------------------------------------------------------------------------------------------------|---------|---------|
| Viewing the instance status of services, configuring parameters, and synchronizing configurations | Yes     | Yes     |
| Viewing cleared alarms and events                                                                 | Yes     | Yes     |
| Viewing the alarm help                                                                            | Yes     | Yes     |
| Setting thresholds for threshold alarms                                                           | No      | Yes     |
| Adding message subscription specifications                                                        | Yes     | No      |
| Managing files                                                                                    | Yes     | No      |
| Managing jobs                                                                                     | Yes     | No      |
| Managing tenants                                                                                  | Yes     | Yes     |
| Managing tags                                                                                     | Yes     | No      |
| Setting permissions (adding and deleting users, user groups, and roles)                           | No      | Yes     |
| Backing up component data and restoring data                                                      | No      | Yes     |
| Logging metadata operation audit                                                                  | No      | Yes     |
| Monitoring resources                                                                              | Yes     | Yes     |

## MRS Cluster Management Objects

There are various fundamental objects in MRS, which are listed in [Table 6-2](#).

**Table 6-2** MRS basic objects

| Object    | Description                                                                     | Example                                                                  |
|-----------|---------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| Component | Function set that can complete specific business.                               | KrbServer and LdapServer                                                 |
| Instance  | Specific instance of a component, which can also be referred to as a component. | KrbServer                                                                |
| Role      | Function entity that forms a complete component, usually called role.           | KrbServer is composed of the KerberosAdmin role and KerberosServer role. |

| Object        | Description                                                                   | Example                                                                                                                       |
|---------------|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Role instance | Specific instance of a service role running on a host.                        | KerberosAdmin that is running on Host2 and KerberosServer that is running on Host3                                            |
| Host          | An ECS running Linux OS.                                                      | Host1 to Host5                                                                                                                |
| Rack          | Physical entity that contains multiple hosts connecting to the same switch.   | Rack1 contains Host1 to Host5.                                                                                                |
| Cluster       | Logical entity that consists of multiple hosts and provides various services. | Cluster names <b>Cluster1</b> consists of five hosts (Host1 to Host5) and provides services such as KrbServer and LdapServer. |

## 6.2 Introduction to MRS Manager

### Manager Overview

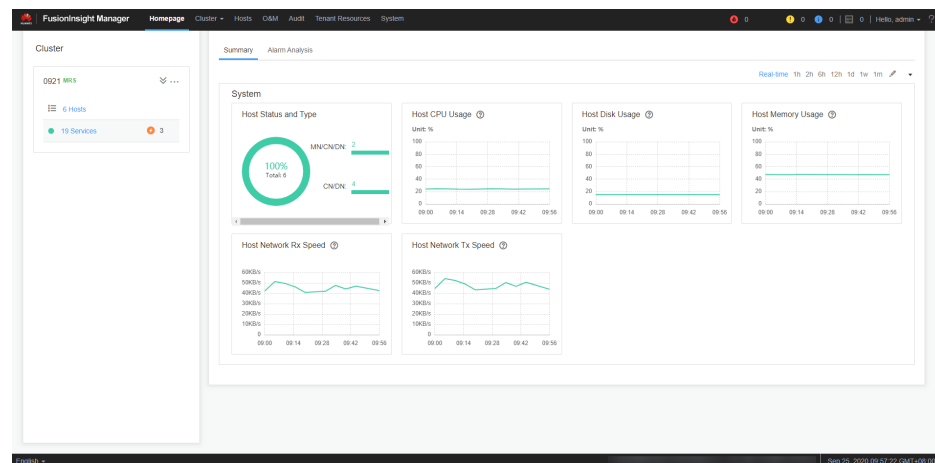
MRS manages and analyzes massive amounts of data and quickly mines valuable data from structured and unstructured data. Open source components have complex structures and therefore they are difficult to install, configure, and manage. MRS Manager is a unified enterprise-level cluster management platform that provides:

- Cluster monitoring: enables you to quickly learn the running status of hosts and services.
- Graphical metric monitoring and customization: enable you to obtain key system information in a timely manner.
- Service property configurations can meet service performance requirements.
- Cluster, service, and role instance operations: allow you to start or stop services and clusters with just a few clicks.
- Rights management and audit: allow you to configure the access control and manage operation logs.

### Manager GUI

MRS Manager provides a unified cluster management platform, facilitating rapid and easy O&M for clusters.

**Figure 6-1** Manager GUI of MRS 3.x

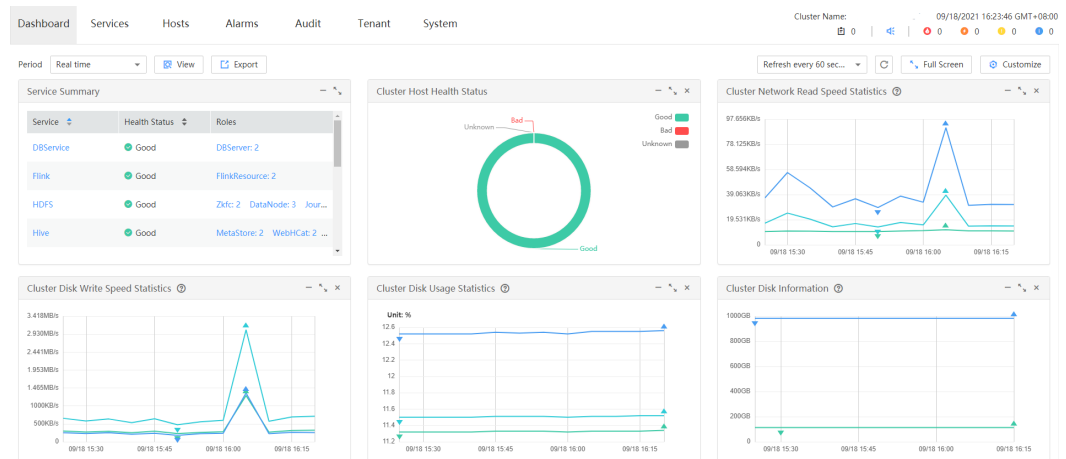


The Manager GUI contains three parts: the operation bar at the top, the display area in the middle, and the taskbar at the bottom. The operation bar has multiple entries, each with its unique function, which are explained in the table provided.

**Table 6-3** Functions of different entries (MRS 3.x)

| Entry            | Function                                                                                                                                                                                                                                                             |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Homepage         | Displays key monitoring metrics of clusters and host statuses in column charts, line charts, and tables. You can customize a dashboard for key monitoring metrics and drag them onto any positions on the GUI. Data on the dashboard can be automatically refreshed. |
| Cluster          | Provides guidance on how to monitor, operate, and configure services in a cluster, helping you manage services in a unified manner.                                                                                                                                  |
| Hosts            | Provides guidance on how to monitor and operate hosts, helping you manage hosts in a unified manner.                                                                                                                                                                 |
| O&M              | Provides guidance on how to query and handle alarms, helping you identify and rectify product faults and potential risks in a timely manner to ensure normal system operations.                                                                                      |
| Audit            | Allows you to query and export audit logs, and view all user activities and operations.                                                                                                                                                                              |
| Tenant Resources | Provides a unified tenant management platform.                                                                                                                                                                                                                       |
| System           | Provides system management settings of FusionInsight Manager, such as user permission settings.                                                                                                                                                                      |

**Figure 6-2** Manager GUI of MRS 2.x



**Table 6-4** Functions of different entries (MRS 2.x)

| Parameter | Function                                                                                                                                                                                                                                                                                                                         |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Dashboard | Displays the status of all services, main monitoring indicators of each service, and host status in charts, such as bar charts, line charts, and tables. You can customize a dashboard for the key monitoring indicators and drag it to any position on the interface. The system dashboard page supports automatic data update. |
| Services  | Provides the service monitoring, operation, and configuration guidance, which helps you manage services in a unified manner.                                                                                                                                                                                                     |
| Hosts     | Provides guidance on how to monitor, operate, and configure hosts, helping you manage hosts in a unified manner.                                                                                                                                                                                                                 |
| Alarms    | Supports alarm query and provides guidance on alarm handling, helping you identify and rectify product faults and potential risks in a timely manner to ensure normal system operation.                                                                                                                                          |
| Audit     | Allows authorized users to query and export audit logs, helping you to view all user activities and operations.                                                                                                                                                                                                                  |
| Tenant    | Provides a unified tenant management platform.                                                                                                                                                                                                                                                                                   |
| System    | Provides monitoring, alarm configuration management, and backup management.                                                                                                                                                                                                                                                      |

## Manager Homepage (MRS 3.x)


On the FusionInsight Manager homepage, you can view the cluster service status, monitoring reports, and alarm statistics and analysis information in the preview area.



- On the right of the homepage, you can view the number of alarms of different severities, number of running tasks, current user, and help information.

**Figure 6-3** Cluster status information




Click  to view the task name, status, progress, start time, and end time of the last 100 operation tasks in **Task Management Center**.

 **NOTE**

For a start, stop, restart, or rolling restart task, you can abort it by clicking the task name in the task list, clicking **Abort**, and then entering the system administrator password in the dialog box that is displayed. An aborted task is no longer executed.




- The taskbar at the bottom of the homepage displays the language options of FusionInsight Manager and the current cluster time and time zone information. You can switch the system language as needed.

**Figure 6-4** Taskbar at the bottom of the homepage




- **Service status preview area**

The list of installed service components in the cluster is displayed on the left of the homepage. You can view the status and alarms of each service.

The  icon on the left of each service name indicates that the service is running properly; the  icon indicates that the current service fails to start; and the  icon indicates that the current service is not started.


You can also check whether alarms have been generated for the service on the right of the service name. If alarms have been generated, the alarm severities and the number of alarms are displayed.

The  icon displayed on the right of the service name indicates that the service configuration has expired.
- **Monitoring status report area**

The chart area is on the right of the homepage, which displays key monitoring metric reports, such as the status of all hosts in the cluster, host CPU usage, and host memory usage. You can customize the monitoring reports displayed in the chart area. For details, see [Viewing MRS Cluster Resource Monitoring Metrics](#).


You can view the data source of a monitoring chart in the lower left corner of the chart. You can zoom in on a monitoring report to view chart values more clearly or close the monitoring report.

Figure 6-5 Monitoring status report



- Alarm analysis  
On the **Alarm Analysis** page, you can view the **Top 20 Alarms** table and **Analysis on Top 3 Alarms** chart. You can click an alarm name in the **Top 20 Alarms** table to view analysis information of this alarm only. Alarm analysis allows you to view top alarms and their occurrence time so you can handle alarms accordingly, improving system stability.

Figure 6-6 Alarm analysis



## Manager Security Functions

You can query and set user rights data through the following Manager modules:

- Role management: Roles can be added, deleted, modified, queried, and assigned with the resource access rights of one or multiple components. For details, see [Managing MRS Cluster Roles](#).
- User group management: User groups can be added, deleted, modified, queried, and bound to roles. For details, see [Managing MRS Cluster User Groups](#).
- User management: Users can be added, deleted, modified, queried, bound to user groups, and assigned with roles. For details, see [Managing MRS Cluster Users](#).
- Tenant management: Tenants can be added, deleted, modified, queried, and bound to component resources. MRS generates a role for each tenant to


facilitate management. If a tenant is assigned with the rights of some resources, its corresponding role also has these rights.

For details, see [Managing MRS Cluster Tenants](#).

## Checking the Manager Version

By checking the Manager version, you can get ready for system upgrades and regular maintenance. This section uses MRS clusters of 3.x or later as an example.

- Using the GUI:

Log in to MRS Manager. On the homepage, click  in the upper right corner and choose **About** from the drop-down list. In the dialog box that is displayed, check the Manager version.

**Figure 6-7** Checking the version number



- Using the CLI:
  - a. Log in to the Manager active OMS node as user **root**.  
To obtain the IP address of the active OMS node, view the host information marked with ★ on the **Hosts** page of Manager.
  - b. Check the version and platform information of Manager.

**su - omm**

**cd \${BIGDATA\_HOME}/om-server/om/sbin/pack**

**./queryManager.sh**

The following information is displayed:

| Version | Package                   | Cputype |
|---------|---------------------------|---------|
| ***     | FusionInsight_Manager_*** | x86_64  |

 **NOTE**

\*\*\* indicates the version number. Replace it with the actual version number.

## 6.3 Accessing MRS Manager


### Scenario

MRS allows you to oversee, adjust, and handle clusters on Manager. Once the cluster is set up, you can access Manager as user **admin**.

Currently, you can access Manager using the following methods:

- **Accessing MRS Manager Using an EIP:** You can bind an EIP to the cluster to access the MRS Manager GUI and open source components managed in the cluster. This method is suggested as it is more user-friendly.
- **Accessing MRS Cluster Manager Using Direct Connect:** Direct Connect is a high-speed, low-latency, stable, and secure dedicated network connection that connects your local data center to an online cloud VPC. It extends online cloud services and existing IT facilities to build a flexible, scalable hybrid computing environment.

You need to ensure that Direct Connect is available, and the connection between the local data center and the online VPC has been established. For details, see [What Is Direct Connect?](#)

You can switch between EIP access and Direct Connect access on the MRS console as follows: Log in to the MRS console, click  next to **MRS Manager** on the **Dashboard** page of the target MRS cluster, and switch between the two access methods on the displayed page.

- **Accessing MRS Manager Through an ECS:** Access Manager through an ECS that is in the same VPC as the MRS cluster. This method is complex and is recommended when the EIP function is not supported.
- **Configuring an SSH Tunnel to Access MRS Manager:** Users and an MRS cluster are in different networks. As a result, an SSH tunnel needs to be created to send users' requests for accessing websites to the MRS cluster and dynamically forward them to the target websites.

## Prerequisites

Make sure the cluster is not in the starting, stopping, stopped, deleting, deleted, or frozen state before accessing MRS Manager.

## Accessing MRS Manager Using an EIP

**Step 1** Log in to the MRS management console.

**Step 2** In the navigation pane, choose **Active Clusters**. Click the target cluster name to access the cluster details page.

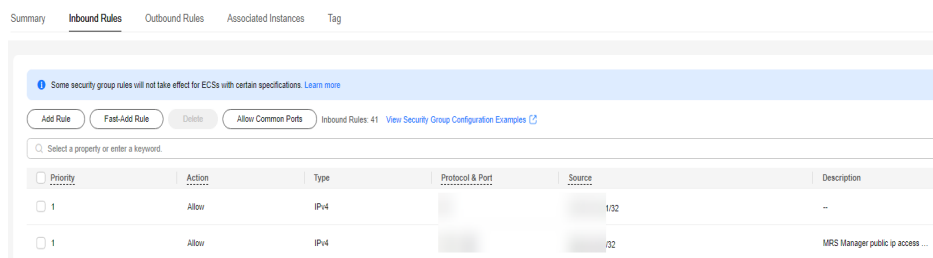
**Step 3** Click **Access Manager** next to **MRS Manager**. In the displayed dialog box, select **EIP** and configure the EIP information.

1. If no EIP is bound when the MRS cluster is created, select an available EIP from the EIP drop-down list. Otherwise, perform the operations in [Step 3.2](#).
  - If no available EIPs are displayed, click **Manage EIP** to create one. An EIP can be bound to only one MRS cluster.
  - To unbind or release an EIP, log in to the **EIPs** page, locate the row containing the target EIP, and click **Unbind** or choose **More > Release** in the **Operation** column.
2. In **Security Group**, select the security group to which the current cluster belongs. The security group is configured during cluster creation or is automatically created by the cluster.

If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

- "MRS Manager public ip access control rule" is automatically added to the **Description** column of the added security group. To view this description, choose **Manage Security Group Rule**, click **Security Group**, and click **Inbound Rules**.

**Figure 6-8** Adding a security group rule to the MRS cluster



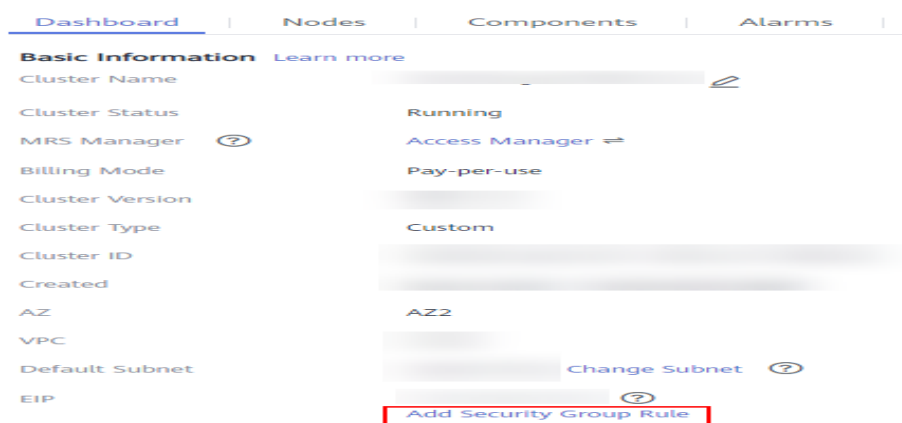
- It is normal that the automatically generated public IP address is different from your local IP address and no action is required.
  - Port **9022** is the Knox port of the MRS cluster. Therefore, you need to enable the permission to access the port to access Manager.
3. Select the information to be confirmed and click **OK**. The Manager login page is displayed.

**Step 4** Enter the default username **admin** and the password set during cluster creation, and click **Log In**. The Manager page is displayed.

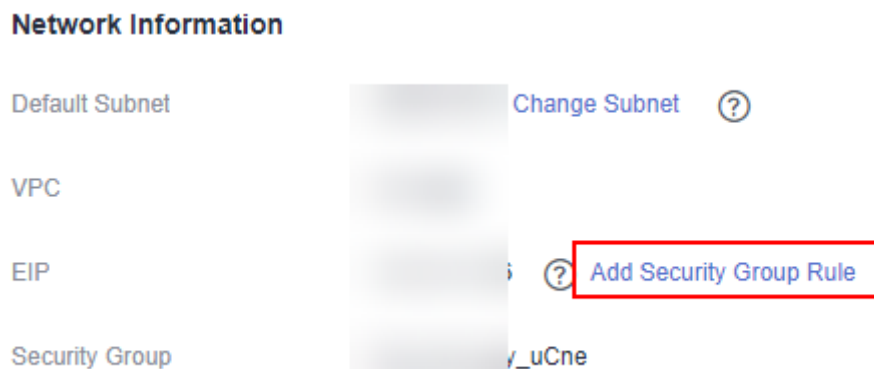
**Step 5** To grant users in other network segments the permission to access Manager, you can modify the security group and add the IP address range for the users to access the public network.

1. Click **Add Security Group Rule** next to **EIP**.

**Figure 6-9** Cluster details



2. On the **Add Security Group Rule** page, add the IP address segment for users to access the public network and select **I confirm that *public network IP/ port* is a trusted public IP address. I understand that using 0.0.0.0/0. poses security risks.** See [Figure 6-10](#).

**Figure 6-10** Adding a security group rule

By default, the IP address segment used for accessing the public network is filled. You can change the IP address segment as required. To view, modify, or delete security group rules, click **Manage Security Group Rule**.

3. Click **OK**.

----End

## Accessing MRS Cluster Manager Using Direct Connect

**Step 1** Log in to the MRS console.

**Step 2** Click the name of the cluster to enter its details page.

**Step 3** On the **Dashboard** page of the cluster details page, click **Access Manager** next to **MRS Manager**.

**Step 4** Set **Access Mode** to **Direct Connect** and confirm that you understand the impact of the operation.

The floating IP address is automatically allocated by MRS to access MRS Manager. Before using Direct Connect to access MRS Manager, ensure that the connection between the local data center and the online VPC has been established.

**Step 5** Click **OK**. The MRS Manager login page is displayed. Enter the username **admin** and the password set during cluster creation.

----End

## Accessing MRS Manager Through an ECS

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, click the name of the specified cluster.

Record the **AZ**, **VPC**, and **Security Group** of the cluster.

**Step 3** On the homepage of the management console, choose **Service List > Elastic Cloud Server** to switch to the ECS management console and create an ECS.

- The **AZ**, **VPC**, and **Security Group** of the ECS must be the same as those of the cluster to be accessed.

- Select a Windows public image. For example, a standard image **Windows Server 2012 R2 Standard 64bit(40GB)**.
- For details about other configuration parameters, see [Purchasing an ECS with Customized Configurations](#).

**NOTE**

If the security group of the ECS is different from **Default Security Group** of the Master node, you can modify the configuration using either of the following methods:

- Change the security group of the ECS to the default security group of the Master node. For details, see [Changing a Security Group](#).
- Add two security group rules to the security groups of the Master and Core nodes to enable the ECS to access the cluster. Set **Protocol** to **TCP**, **Ports** of the two security group rules to **28443** and **20009**, respectively. For details, see [Creating a Security Group](#).

If "Failed to add security group rules." is displayed, check whether the security group quota is sufficient. If more quotas are needed, increase the quotas or delete security group rules that are no longer used.

**Step 4** On the EIP console, apply for an EIP and bind it to the ECS.

For details, see [Assigning an EIP and Binding It to an ECS](#).

**Step 5** Log in to the ECS.


The Windows system account, password, EIP, and security group rules are required for logging in to the ECS. For details, see [Logging In to a Windows ECS](#).

**Step 6** On the Windows remote desktop, use your browser to access Manager.

The Manager access address is in the **https://OMS floating IP address:28443/web** format. Enter the name and password of the cluster user, for example, user **admin**.

**NOTE**

- To obtain the floating IP address of OMS, log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of OMS. Record the value of **inet**. If the floating IP address of OMS cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one Master node, query and record the cluster manager IP address of the Master node.
- If you access Manager with other cluster usernames, change the password upon your first access. The new password must meet the requirements of the current password complexity policies. For details, contact the administrator.
- By default, a user is locked after inputting an incorrect password five consecutive times. The user is automatically unlocked after 5 minutes.

**Step 7** Log out of FusionInsight Manager. To log out of Manager, move the cursor to  in the upper right corner and click **Log Out**.

----End

## Configuring an SSH Tunnel to Access MRS Manager

Users and an MRS cluster are in different networks. As a result, an SSH tunnel needs to be created to send users' requests for accessing websites to the MRS cluster and dynamically forward them to the target websites.

The MAC system does not support this function. For details about how to access MRS, see [Accessing MRS Manager Using an EIP](#).

Make sure the following prerequisites are met before proceeding with the operation:

- You have prepared an SSH client for creating the SSH tunnel, for example, the Git open source SSH client. You have downloaded and installed the client.
- You have created a cluster and prepared a key file in PEM format or obtained the password used during cluster creation.
- Users can access the Internet on the local PC.

**Step 1** Log in to the MRS console and click **Active Clusters**.

**Step 2** Click the specified MRS cluster name.

Record the security group of the cluster.

**Step 3** Add an inbound rule to the security group of the master node to allow data access to the IP address of the MRS cluster through port 22.

For details, see [Adding a Security Group Rule](#).

**Step 4** Query the primary management node of the cluster by referring to [Checking MRS Active/Standby Management Nodes](#).

**Step 5** Bind an elastic IP address to the primary management node of the cluster.

For details, see [Assigning an EIP and Binding It to an ECS](#).

**Step 6** Start Git Bash locally and run the following command to log in to the active management node of the cluster:

```
ssh root@EIP address
```

Alternatively, run the following command:

```
ssh -i Key file path root@EIP address
```

**Step 7** View data forwarding configurations.

```
cat /etc/sysctl.conf | grep net.ipv4.ip_forward
```

- If **net.ipv4.ip\_forward=1** is displayed, the forwarding function has been configured. Go to [Step 9](#).
- If **net.ipv4.ip\_forward=0** is displayed, the forwarding function has not been configured. Go to [Step 8](#).
- If **net.ipv4.ip\_forward** fails to be queried, this parameter has not been configured. Run the following command and then go to [Step 9](#):  
echo "net.ipv4.ip\_forward = 1" >> /etc/sysctl.conf

**Step 8** Modify forwarding configurations on the node.

1. Switch to user **root**.

```
sudo su - root
```

2. Modify forwarding configurations.

```
echo 1 > /proc/sys/net/ipv4/ip_forward
```

```
sed -i "s/net.ipv4.ip_forward=0/net.ipv4.ip_forward = 1/g" /etc/sysctl.conf
```



```
sysctl -w net.ipv4.ip_forward=1
```

3. Modify the `sshd` configuration file.

```
vi /etc/ssh/sshd_config
```

Press **I** to enter the edit mode. Locate **AllowTcpForwarding** and **GatewayPorts** and delete comment tags. Modify them as follows. Save the changes and exit.

```
AllowTcpForwarding yes
GatewayPorts yes
```

4. Restart the `sshd` service.

```
service sshd restart
```

- Step 9** View the floating IP address.

```
ifconfig
```

In the command output, **eth0:FI\_HUE** indicates the floating IP address of Hue, and **eth0:wsom** indicates the floating IP address of Manager. Record the value of **inet**.

Run the **exit** command to exit.

- Step 10** Run the following command on the local host to create an SSH tunnel that supports dynamic port forwarding:

```
ssh -i Path of the key file -v -ND Local port root@EIP address
```

Alternatively, run the following command:

```
ssh -v -ND Local port root@EIP address
```

Enter the password for creating the cluster as prompted.

In the command, set **Local port** to the user's local port that is not occupied. Port **8157** is recommended.

After the SSH tunnel is created, use **-D** to enable the dynamic port forwarding function. By default, the dynamic port forwarding function enables a SOCKS proxy process and monitors the user's local port. Port data will be forwarded to the primary management node using the SSH tunnel.

- Step 11** Configure the browser proxy.

1. Go to the Google Chrome client installation directory on the local PC.
2. Press **Shift** and right-click the blank area, choose **Open Command Window Here** and enter the following command:

```
chrome --proxy-server="socks5://localhost:8157" --host-resolver-rules="MAP * 0.0.0.0 , EXCLUDE localhost" --user-data-dir=c:/tmp/path --proxy-bypass-list="*google*.com,*gstatic.com,*gvt*.com,*:80"
```

 **NOTE**

- In the preceding command, **8157** is the local proxy port configured in [Step 10](#).
- If the local OS is Windows 10, start the Windows OS, click **Start** and enter **cmd**. In the displayed CLI, run the command in [Step 11.2](#). If this method fails, click **Start**, enter the command in the search box, and run the command in [Step 11.2](#).

- Step 12** In the address box of the browser, enter the address for accessing Manager.

The Manager access address is in the **https://Manager floating IP address:28443/web** format.

The username and password of the MRS cluster need to be entered for accessing clusters with Kerberos authentication enabled, for example, user **admin**. They are not required for accessing normal clusters with Kerberos authentication disabled.

For the first access, add the site to the trusted site list as prompted to continue to open the page.

**Step 13** When logging out of Manager, terminate and close the SSH tunnel.

----End

## 6.4 Managing an MRS Cluster

### 6.4.1 Viewing Basic Information About an MRS Cluster

You can monitor and manage the clusters you have created. On the **Active Clusters** page. Click the name of a cluster to go to the cluster details page. On the displayed page, view the basic configuration, network, and node information of the cluster.

#### NOTE

On the MRS console, operations performed on an ECS cluster are basically the same as those performed on a BMS cluster. This document describes operations on an ECS cluster. If operations on the two clusters differ, the operations will be described separately.


### Viewing Basic Information About an MRS Cluster

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the cluster details page, click **Dashboard** to view the basic information of the cluster.

**Table 6-5** MRS cluster basic information

| Parameter       | Description                                                                                                                                                                                                                                                                                                                                  |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Name    | The name of a cluster. Configure this parameter when creating a cluster. Click  to change the cluster name. Only the cluster name displayed on the MRS management console is changed, while the cluster name on MRS Manager is not changed synchronously. |
| Cluster Status  | The cluster status. For details, see <a href="#">Table 6-10</a> .                                                                                                                                                                                                                                                                            |
| Cluster Version | MRS version information.                                                                                                                                                                                                                                                                                                                     |

| Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Type            | There are three types of clusters: <ul style="list-style-type: none"><li>● <b>Analysis Cluster</b>: is used for offline data analysis and provides Hadoop components.</li><li>● <b>Streaming Cluster</b>: is used for streaming tasks and provides stream processing components.</li><li>● <b>Hybrid Cluster</b>: is used for both offline data analysis and streaming processing and provides Hadoop components and streaming processing components.</li><li>● <b>Custom</b>: An MRS cluster with all custom components. MRS 3.x or later supports this type.</li></ul> |
| Cluster ID              | Unique identifier of a cluster, which is automatically assigned when a cluster is created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Created                 | Time when a cluster is created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| AZ                      | Availability zone (AZ) in the region of a cluster, which is set when a cluster is created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Kerberos Authentication | Whether to enable Kerberos authentication when logging in to Manager.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Enterprise Project      | The enterprise project to which a cluster belongs. Only on the <b>Active Clusters</b> page, you can click the name of an enterprise project to go to its <b>Enterprise Project Management</b> page.                                                                                                                                                                                                                                                                                                                                                                      |

**Table 6-6** MRS cluster network information

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Default Subnet | The subnet selected during cluster creation.<br>If the subnet IP addresses are insufficient, click <b>Change Subnet</b> to switch to another subnet in the same VPC of the current cluster to obtain more available subnet IP addresses. Changing a subnet does not affect the IP addresses and subnets of existing nodes.<br>A subnet provides dedicated network resources that are isolated from other networks, improving network security. |
| VPC            | VPC selected during cluster creation.<br>A VPC is a secure, isolated, and logical network environment.                                                                                                                                                                                                                                                                                                                                         |

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                              |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EIP            | <p>After binding an EIP to an MRS cluster, you can use the EIP to access the Manager web UI of the cluster. If you do not need the EIP, click <b>Unbind</b> to unbind the EIP from the cluster. The Manager will not be accessible through the unbound EIP.</p> <p><b>NOTE</b><br/>If you unbind the EIP from a cluster, other users may fail to access the Manager of that cluster.</p> |
| Security Group | The security group name of the cluster.                                                                                                                                                                                                                                                                                                                                                  |

**Table 6-7** MRS cluster O&M management

| Parameter              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MRS Manager            | Manager entry. For details, see <a href="#">Accessing MRS Manager</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| IAM User Sync          | <p>IAM user information (including federated users) can be synchronized to an MRS cluster for cluster management. For details, see <a href="#">Synchronizing IAM Users to MRS</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The <b>Components</b>, <b>Tenants</b>, and <b>Backups &amp; Restorations</b> pages on the cluster details page can be used only after users are synchronized. After an MRS 3.x cluster is synchronized, you can use the <b>Components</b> function.</li> <li>For a federated user, only information about the user that logged in the system can be synchronized.</li> </ul> |
| Data Connection        | Click <b>Manage</b> to view the data connection type associated with the cluster. For details, see <a href="#">Creating a Data Connection</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Agency                 | <p>Click <b>Manage Agency</b> to bind or modify an agency for the cluster.</p> <p>An agency allows ECS or BMS to manage MRS resources. You can configure an agency of the ECS type to automatically obtain the AK/SK to access OBS. For details, see <a href="#">Interconnecting an MRS Cluster with OBS Using an IAM Agency</a>.</p> <p>The <b>MRS_ECS_DEFAULT_AGENCY</b> agency has the OBSOperateAccess permission of OBS and the CESFullAccess (for users who have enabled fine-grained policies), CES Administrator, and KMS Administrator permissions in the region where the cluster is located.</p>                   |
| OBS Permission Control | Click <b>Manage</b> and modify the mapping between MRS users and OBS permissions. For details, see <a href="#">Configuring Fine-Grained OBS Access Permissions for MRS Cluster Users</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                    |

| Parameter             | Description                                                                                                                                                                                                                                                                                                    |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Logging               | Used to collect logs about cluster creation and scaling failures.                                                                                                                                                                                                                                              |
| Secure Communications | Used to display the security authorization status. You can enable or disable security authorization. Disabling security authorization brings high risks. Exercise caution when performing this operation. For details, see <a href="#">Configuring Secure Communication Authorization for an MRS Cluster</a> . |

**Table 6-8** MRS cluster billing information

| Parameter              | Description                                                                                                                                          |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Billing Mode           | Billing mode of a cluster. Currently, <b>Pay-per-use</b> and <b>Yearly/Monthly</b> are supported.                                                    |
| Last Transaction Order | Order number for purchasing a cluster. This parameter is available only when <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> .                   |
| Created                | Time when the cluster is created. This parameter is available only when <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> .                        |
| Expired                | Cluster expiration time. This parameter is available only when <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> .                                 |
| Upon Expiration        | The cluster will enter the grace period upon expiration. This parameter is available only when <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> . |

----End

## 6.4.2 Checking the Running Status of an MRS Cluster

MRS allows you to buy multiple clusters. The cluster quantity is subject to that of ECSs. You can view the running status of all MRS clusters on the console.

### Checking the Running Status of an MRS Cluster

**Step 1** Log in to the MRS console.



**Step 2** Click **Active Clusters**.

By default, clusters in the cluster list are sorted by creation time. For details about the cluster list parameters, see [Table 6-9](#).

- **Active Clusters**: contain all clusters except the clusters in the **Failed** and **Deleted** states.
- **Cluster History**: contains the clusters in the **Deleted** states. Only clusters deleted within the last six months are displayed. If you want to view clusters deleted six months ago, contact Huawei Cloud technical support.

- **Failed Tasks:** Click  to view the failed cluster creation tasks.

**Table 6-9** Parameters in the active cluster list

| Parameter       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name/ID         | <p>Cluster name, which is set when a cluster is created. Unique identifier of a cluster, which is automatically assigned when a cluster is created.</p> <ul style="list-style-type: none"> <li>•  : Change the cluster name.</li> <li>•  : Copy the cluster ID.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Cluster Version | Cluster version.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Cluster Type    | The type of a cluster you want to create.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Nodes           | Number of nodes that can be deployed in a cluster. This parameter is set when a cluster is created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Status          | <p>Cluster running state and state change information. For details, see <a href="#">Table 6-10</a>.</p> <p>The cluster creation progress includes:</p> <ul style="list-style-type: none"> <li>• Verifying cluster parameters</li> <li>• Applying for cluster resources</li> <li>• Creating VMs</li> <li>• Initializing VMs</li> <li>• Installing MRS Manager</li> <li>• Deploying the cluster</li> <li>• Cluster installation failed</li> </ul> <p>The cluster scale-out progress includes:</p> <ul style="list-style-type: none"> <li>• Preparing for scale-out</li> <li>• Creating VMs</li> <li>• Initializing VMs</li> <li>• Adding nodes to the cluster</li> <li>• Scale-out failed</li> </ul> <p>The cluster scale-in progress includes:</p> <ul style="list-style-type: none"> <li>• Preparing for scale-in</li> <li>• Decommissioning instance</li> <li>• Deleting VMs</li> <li>• Deleting nodes from the cluster</li> <li>• Scale-in failed</li> </ul> <p>The system will display causes of cluster installation, scale-out, and scale-in failures. For details, see <a href="#">Table 3-7</a>.</p> |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Billing Mode       | <p>Currently, the commercial version of MRS is charged based on ECSs in a cluster.</p> <ul style="list-style-type: none"><li>• <b>Yearly/Monthly:</b> The duration ranges from one month to one year. The minimum cluster duration is 1 month and the maximum available cluster duration is 1 year.</li><li>• <b>Pay-per-use:</b> Nodes are charged by actual duration of use, with a billing cycle of one hour.</li></ul> <p>The billing start time, that is, the time when the cluster node is successfully created is displayed under <b>Billing Mode</b>.</p> |
| Created            | The time when a cluster node is successfully created. This parameter is displayed only on the <b>Cluster History</b> page.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Deleted            | Time when a cluster node billing stops and the cluster node begins to be deleted. This parameter is valid only for historical clusters displayed on the <b>Cluster History</b> page.                                                                                                                                                                                                                                                                                                                                                                              |
| AZ                 | Availability zone (AZ) in the region of a cluster, which is set when a cluster is created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Enterprise Project | Enterprise project to which a cluster belongs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Table 6-10** MRS cluster running state

| Status      | Description                                                                                                                                                                                                                                             |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Starting    | When a cluster is being created, it is in the <b>Starting</b> state.                                                                                                                                                                                    |
| Running     | If a cluster is created successfully and functioning properly, it is in the <b>Running</b> state.                                                                                                                                                       |
| Scaling out | If the Master, Core, or Task node in a cluster is being added, the cluster is in the <b>Scaling out</b> state.                                                                                                                                          |
| Scaling in  | If a cluster node is being deleted, the cluster node is in the <b>Scaling in</b> state. This state shows when you scale in or elastically scale in a cluster node, decommission a node in a yearly/monthly cluster, change the OS, or reinstall the OS. |
| Abnormal    | If some components in a cluster are abnormal, the cluster is in the <b>Abnormal</b> state.                                                                                                                                                              |
| Deleting    | When you delete a pay-per-use cluster node, the cluster state changes to <b>Deleting</b> . This state is shown after you click <b>Delete</b> and confirm the deletion.                                                                                  |

| Status            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Frozen            | If the grace period of a yearly/monthly resource expires but the resource is not renewed, or if the fee of a pay-per-use resource fails to be deducted and you have not topped up your account before the grace period expires, the system freezes the resource in the <b>Frozen</b> state.<br><b>NOTE</b><br>A frozen cluster is unavailable and its all ECSs are shut down. After being unfrozen, the cluster returns to the <b>Running</b> state. If no renewal fee is paid, the cluster will be deleted after a specified period (called freeze period) and the cluster status will be changed to <b>Deleted</b> . |
| Restoring node... | If a faulty node in the cluster is being recovered, its state is <b>Restoring node....</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

----End

### 6.4.3 Starting and Stopping an MRS Cluster

You can stop an MRS cluster using the MRS console or Manager when it is no longer needed or has been fixed due to a fault. Once stopped, the cluster's components will no longer serve external systems.

You can also restart a stopped cluster.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

#### Starting and Stopping a Cluster on the Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.


**Step 3** On the cluster details page, choose **Management Operations** > **Start All Components** or **Stop All Components** in the upper right corner to perform the required operation.

----End

#### Starting and Stopping a Cluster Using Manager

- MRS 3.x and later versions:
  - a. Log in to Manager.



- b. Click  next to the target cluster and select **Stop**. Enter the password of the cluster administrator. In the **Stop Cluster** dialog box that is displayed, click **OK**. Wait until the cluster is stopped.
- MRS 2.x and earlier versions:
  - a. Log in to Manager and click **Services**.
  - b. In the upper part of the service list, choose **More > Start Cluster** or **Stop Cluster** accordingly.

## 6.4.4 Restarting an MRS Cluster

To apply configuration changes to a big data component, you must restart it. However, using the common restart mode will restart all services or instances at once, which can cause service interruption.

To ensure that services are not affected during service restart, you can restart services or instances in batches by rolling restart. For instances in active/standby mode, a standby instance is restarted first and then an active instance is restarted.

A rolling restart takes a longer time and may affect service throughput and performance.

### NOTE

- Performing a rolling restart during off-peak hours is recommended.
- Not all components support rolling restart. When performing a rolling restart of the cluster, components that do not support it will be restarted in common restart mode, which may cause service interruptions. For details, see [Component Restart Reference Information](#).
- Configurations that must take effect immediately, for example, server port configurations, should be restarted in normal mode.


## Restarting a Cluster on the Console

**Step 1** Log in to the MRS console.

**Step 2** Choose **Active Clusters** and click a cluster name to go to the cluster details page.

**Step 3** In the upper right corner of the page, choose **Management Operations > Perform Rolling Cluster Restart**.

**Figure 6-11** Performing a rolling restart of a cluster (use MRS 1.9.2 as an example)



**Step 4** The **Rolling-restart Cluster** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the target cluster.

**Step 5** After the rolling restart task is complete, click **Finish**.

----End

## Restarting a Cluster on Manager

- MRS 3.x or later:
  - a. Log in to FusionInsight Manager.
  - b. Choose **Cluster > Dashboard**. In the upper right corner, click **More > Restart**.


### NOTE

- For MRS 3.3.0 or later, the **Cluster > Dashboard** page has been removed from Manager. You can choose **More** in the upper right corner of the **Homepage** to access cluster maintenance and management functions.
  - When restarting a cluster, you can choose between restart or rolling restart. Rolling restart takes more time but has less impact on services.
- c. In the displayed dialog box, enter the password of the current login user and click **OK**.
  - d. If you select rolling restart, adjust related parameters based on the site requirements.

Figure 6-12 Rolling restart

### Rolling-restart Cluster

**Are you sure you want to perform a rolling restart of the MRS Cluster?**


 A rolling restart can minimize the impact on services, but takes a longer time than a standard restart.

The following services do not support the rolling restart and can be restarted only in common mode:


- Tez
- Oozie
- Loader
- Hue
- Flink


Rolling Restart Options


Restart only instances with expired configurations in the cluster [View Instances](#)


Enable rack strategy 

^ Advanced Options

Data Nodes to Be Batch Restarted: 

Batch Interval:   s

Decommissioning Timeout Interval:   s

Batch Fault Tolerance Threshold: 

**Table 6-11** Rolling restart parameters

| Parameter                                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Restart only instances with expired configurations in the cluster | Whether to restart only the modified instances in a cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Enable rack strategy                                              | Whether to enable the concurrent rack rolling restart strategy. This parameter takes effect only for roles that meet the rack rolling restart strategy. (The roles support rack awareness, and instances of the roles belong to two or more racks.)<br><b>NOTE</b><br>This parameter can be set only when a rolling restart is performed on HDFS or YARN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Data Nodes to Be Batch Restarted                                  | Number of instances that are restarted in each batch when the batch rolling restart strategy is used. The default value is <b>1</b> .<br><b>NOTE</b> <ul style="list-style-type: none"> <li>This parameter is valid only when the batch rolling restart strategy is used and the instance type is DataNode.</li> <li>This parameter is invalid when the rack strategy is enabled. In this case, the cluster uses the maximum number of instances (<b>20</b> by default) configured in the rack strategy as the maximum number of instances that are concurrently restarted in a rack.</li> <li>This parameter is configurable only when a rolling restart is performed on HDFS, HBase, YARN, Kafka, Storm, or Flume.</li> <li>This parameter for the RegionServer of HBase cannot be manually configured. Instead, it is automatically adjusted based on the number of RegionServer nodes. Specifically, if the number of RegionServer nodes is less than 30, the parameter value is <b>1</b>. If the number is greater than or equal to 30 and less than 300, the parameter value is <b>2</b>. If the number is greater than or equal to 300, the parameter value is 1% of the number (rounded-down).</li> </ul> |
| Batch Interval                                                    | Interval between two batches of instances to be roll-restarted. The default value is <b>0</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Decommissioning Timeout Interval | <p>Decommissioning interval for role instances during a rolling restart. The default value is <b>1800s</b>.</p> <p>Some roles (such as HiveServer and JDBCServer) stop providing services before the rolling restart. Stopped instances cannot be connected to new clients. Existing connections will be completed after a period of time. An appropriate timeout interval can ensure service continuity.</p> <p><b>NOTE</b><br/>This parameter is configurable only when a rolling restart is performed on Hive or Spark2x.</p> |
| Batch Fault Tolerance Threshold  | <p>Tolerance times when the rolling restart of instances fails to be batch executed. The default value is <b>0</b>, which indicates that the rolling restart task ends after any batch of instances fails to restart.</p>                                                                                                                                                                                                                                                                                                        |

 **NOTE**

Advanced parameters, such as **Data Nodes to Be Batch Restarted**, **Batch Interval**, and **Batch Fault Tolerance Threshold**, should be properly configured based on site requirements. Otherwise, services may be interrupted or cluster performance may be severely affected.

Examples:

- If **Data Nodes to Be Batch Restarted** is set to an unnecessarily large value, a large number of instances are restarted concurrently. As a result, services are interrupted or cluster performance is severely affected due to too few working instances.
  - If **Batch Fault Tolerance Threshold** is too large, services will be interrupted because a next batch of instances will be restarted after a batch of instances fails to restart.
- e. Click **OK**.
- MRS 2.x and earlier versions:
    - a. On MRS Manager, click **Services**.
    - b. Choose **More > Perform Rolling Service Restart**.
    - c. After you enter the administrator password, the **Perform Rolling Cluster Restart** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the cluster.
    - d. After the rolling restart task is complete, click **Finish**.

## 6.4.5 Exporting MRS Cluster Configuration Parameters

Administrators can export key server configuration parameters of MRS cluster components for quick check or backup.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).


## Exporting Cluster Configurations on the Console

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** Choose **Configuration > Export Cluster Configuration** in the upper right corner of the page to export cluster configurations to the local PC.

### NOTE

This operation applies to MRS 2.x and earlier versions.

**Figure 6-13** Exporting cluster configurations



----End

## Exporting Cluster Configurations Using Manager

- Step 1** Log in to MRS Manager.
- Step 2** Choose **Cluster > Dashboard**.
- Step 3** Choose **More > Export Configurations** to export the configurations of all services in the cluster in batches.

----End

### 6.4.6 Synchronizing the MRS Cluster Configuration

If a new configuration needs to be delivered to all services in the cluster, or **Configuration Status** of multiple services changes to **Expired** or **Failed** after a configuration is modified, the configuration parameters of these services are not synchronized and do not take effect. In this case, synchronize the configurations and restart related service instances for the cluster so that the new parameters take effect for all services.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

## Impact on the System

- After the cluster configurations are synchronized, you need to restart the services whose configurations have expired. During the restart, the corresponding services are unavailable.
- The instances whose configuration has expired are unavailable during restart.

## Synchronizing Cluster Configurations on the Console

**Step 1** Log in to the MRS console.


**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** choose **Configure** > **Synchronize Configuration** in the upper right corner.

### NOTE

This operation applies to MRS 2.x and earlier versions.

**Figure 6-14** Synchronizing configurations (using MRS 1.9.2 as an example)



**Step 4** In the displayed dialog box, select "Restart services and instances whose configuration have expired" and click **OK** to restart the service whose configurations have expired.

When **Operation successful** is displayed, click **Finish**. The cluster is started.

----End

## Synchronizing Cluster Configurations Using Manager

**Step 1** Log in to MRS Manager.

**Step 2** Choose **Cluster** > **Dashboard**.

**Step 3** On this page, choose **More** > **Synchronize Configuration**.

**Step 4** In the dialog box that is displayed, click **OK**.

**Step 5** On the cluster dashboard page, choose **More** > **Restart Configuration-Expired Instances**.

**Step 6** In the dialog box that is displayed, enter the password of the current login user and click **OK**.

**Step 7** In the displayed dialog box, click **OK**.

You can click **View Instance** to open the list of all expired instances and confirm that the instances have been restarted.

----End

## 6.4.7 Transforming a Pay-per-Use MRS Cluster to a Yearly/Monthly Cluster

This section describes how to change the billing mode of a cluster from **Pay-per-use** to **Yearly/Monthly**.

This operation can be performed only when the cluster status is **Running** or **Stopping**.

### Procedure

**Step 1** Log in to the MRS console.

**Step 2** In the navigation pane on the left, choose **Active Clusters**.

**Step 3** In the **Operation** column corresponding to the cluster for which you want to change the billing mode, click **Change to Yearly/Monthly**.

**Step 4** If you are sure you want to change the billing mode, click **Yes**.

**Step 5** On the **Change Subscription** page that is displayed, choose how often you would like to renew and click **Pay**.

After the order is submitted, the cluster status changes from **Running** to **Changing to Yearly/Monthly**.

After the order is paid successfully, the cluster billing mode starts changing to **Yearly/Monthly**. After the billing mode is successfully changed, the cluster status is **Running**.

#### NOTE

After the billing mode is changed to yearly/monthly, Task nodes in a cluster are still billed in pay-per-use mode. During the change, the configured AS rules do not trigger scaling actions. Change the billing mode at an appropriate time to avoid any adverse impact on your services.

----End

## 6.4.8 Deleting an MRS Cluster

To delete or unsubscribe from a cluster, wait until data analysis and storage are finished, or if the cluster is not functioning properly and cannot provide services. If an MRS cluster fails to be deployed, the cluster is automatically deleted or unsubscribed.

- You can delete a pay-per-use cluster if you no longer need it after completing a job. The deleted or unsubscribed cluster is no longer billed.
- You cannot delete a cluster that is billed yearly/monthly if you no longer need it after completing a job. You can unsubscribe from this cluster.

After a cluster is unsubscribed, resources and data will be deleted and cannot be restored. Ensure that the data is backed up before unsubscribing from the cluster. For details about the unsubscription rules, see [Conditional Unsubscription](#).

 **NOTE**

After deleting a component or cluster connected to OBS (including scenarios with storage and compute decoupled or where hot and cold datasets are stored separately), you must manually delete the service data on OBS.

## Deleting an MRS Pay-per-Use Cluster

- Step 1** Log in to the MRS management console.
- Step 2** In the navigation pane on the left, choose **Active Clusters**.
- Step 3** In the cluster list, locate the row containing the cluster to be deleted, and click **Delete** in the **Operation** column. In the **Delete Cluster** dialog box, enter **DELETE** in the confirmation text box and click **OK**.

The cluster status changes from **Running** to **Deleting**, and finally to **Deleted**. You can view the deleted cluster in **Cluster History**. The deleted cluster is no longer billed.

----End

## Unsubscribing from a Yearly/Monthly MRS Cluster

- Step 1** Log in to the MRS management console.
- Step 2** In the navigation pane on the left, choose **Active Clusters**.
- Step 3** In the **Operation** column of the cluster from which you want to unsubscribe, click **Unsubscribe**.
- Step 4** On the **Unsubscribe** page, confirm the cluster information, select reasons for unsubscription, and confirm the unsubscription amount and related fees.
- Step 5** Click **Confirm**.
- Step 6** Confirm the unsubscription information and click **Yes** to submit the unsubscription application.

After the unsubscription application is submitted, the cluster status changes from **Running** to **Deleting**. After the cluster is deleted, the cluster status changes to **Deleted** and is displayed in **Cluster History**.

----End

### 6.4.9 Changing the VPC Subnet of an MRS Cluster

If the current subnet does not have sufficient IP addresses, you can change to another subnet in the same VPC of the current cluster to obtain more available subnet IP addresses. Changing a subnet does not affect the IP addresses or subnets of existing nodes.

For details about how to configure network ACL outbound rules, see [How Do I Configure a Network ACL Outbound Rule?](#)



## Changing a Subnet When No Network ACL Is Associated

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** In the **Network Information** area, click **Change Subnet** on the right of **Default Subnet**.
- Step 4** Select the target subnet and click **OK**.


If no subnet is available, click **Create Subnet** to create a subnet first.

----End

## Changing a Subnet When a Network ACL Is Associated

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** In the **Network Information** area, view **VPC**.
- Step 4** Log in to the VPC console. In the navigation pane on the left, choose **Virtual Private Cloud** and obtain the IPv4 CIDR block corresponding to the VPC obtained in [Step 3](#).

**Figure 6-15** Obtaining the IPv4 CIDR block



- Step 5** Choose **Access Control > Network ACLs** and click the name of the network ACL that is associated with the default and new subnets.

**NOTE**

If both the default and new subnets are associated with a network ACL, add inbound rules to the network ACL by referring to [Step 6](#) to [Step 8](#).

- Step 6** On the **Inbound Rules** page, choose **More > Insert Rule Above** in the **Operation** column.
- Step 7** Add a network ACL rule. Set **Action** to **Allow**, **Source** to the VPC IPv4 CIDR block obtained in [Step 4](#), and retain the default values for other parameters.
- Step 8** Click **OK**.

**NOTE**

If you do not want to allow access from all IPv4 CIDR blocks of the VPC, add the IPv4 CIDR blocks of the default and new subnets by performing [Step 9](#) to [Step 13](#). If you have added rules for the IPv4 CIDR block of the VPC, skip these steps.

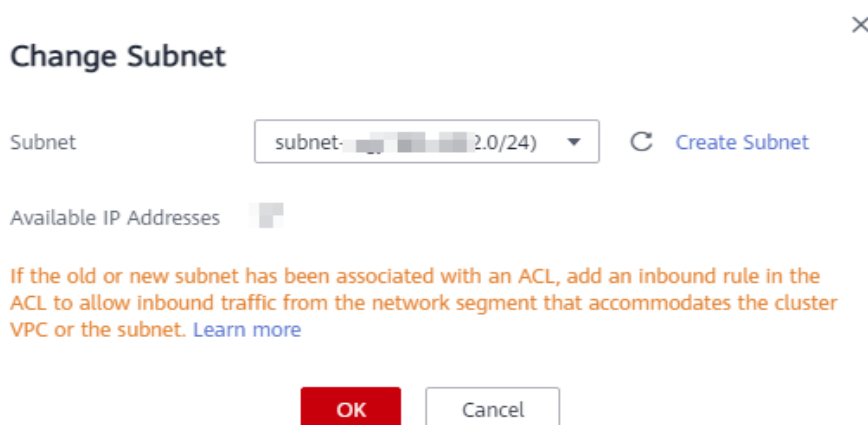
- Step 9** Log in to the MRS console.
- Step 10** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 11** In the **Network Information** area, click **Change Subnet** on the right of **Default Subnet**.
- Step 12** Obtain the IPv4 CIDR blocks of the default and new subnets.

**NOTICE**

In this case, you do not need to click **OK** displayed in the **Change Subnet** dialog box. Otherwise, the default subnet will be updated to the new subnet, thereby making it difficult to query the IPv4 CIDR block of the default subnet. Exercise caution when performing this operation.

- Step 13** Add the default subnet and the new subnet to the inbound rules of the network ACL by referring to [Step 5](#) to [Step 8](#).
- Step 14** Log in to the MRS console.
- Step 15** Click the cluster name to go to its details page.
- Step 16** In the **Network Information** area, click **Change Subnet** on the right of **Default Subnet**.
- Step 17** Select the target subnet and click **OK**.

**Figure 6-16** Selecting the target subnet



----End

## How Do I Configure a Network ACL Outbound Rule?

- Method 1  
Allow all outbound traffic. This method ensures that clusters can be created and used properly.

**Figure 6-17** Allowing all outbound traffic

| Status  | Type | Action | Protocol | Source    | Source Port Range | Destination | Destination Port Range |
|---------|------|--------|----------|-----------|-------------------|-------------|------------------------|
| Enabled | IPv4 | Allow  | All      | 0.0.0.0/0 | All               | 0.0.0.0/0   | All                    |

- Method 2  
Allow the mandatory outbound rules that can ensure the successful creation of clusters. You are not advised to use this method because created clusters may not run properly due to absent outbound rules. If the preceding problem occurs, contact O&M personnel.  
Similar to the example provided in method 1, set **Action** to **Allow** and add the outbound rules whose destinations are the address with **Secure Communications** enabled, NTP server address, OBS server address, OpenStack address, and DNS server address, respectively.

### 6.4.10 Replacing the NTP Server for an MRS Cluster

If no NTP server is configured or the configured NTP server is no longer used, you can specify a new NTP server for the MRS cluster or replace the NTP server with a new one to enable the cluster to synchronize time with the new NTP clock source.

 **NOTE**

This section applies only to MRS 3.x or later.

#### Prerequisites

- You have prepared a new NTP server and obtained its IP address, and have configured the network between the cluster and the new NTP server.
- Ensure that the NTP service status of the server is normal. Otherwise, the operations in this section will fail.

#### Impact on the System

- Replacing the NTP server is a high-risk operation and may result in time change in the cluster.
- If the time difference between the NTP server and the cluster is greater than 150s before the NTP server replacement, you need to stop the cluster first to prevent data loss. Services are unavailable when the cluster is stopped.
- If the time difference between the NTP server and the cluster exceeds 15 minutes, the cluster will be unable to access OBS.
- If your clusters use Kerberos authentication and the time difference between the NTP server and the cluster exceeds 5 minutes, authentication will not work.

## Modifying the NTP Server of an MRS Cluster

**Step 1** Log in to FusionInsight Manager and check whether there are uncleared alarms.

- If yes, clear the alarm. After the alarm is cleared, go to [Step 2](#).
- If no, go to [Step 2](#).

**Step 2** Log in to the active and standby management nodes as user **omm**.

**Step 3** Run the following command on the active management node to check the management plane gateway:

```
cat ${BIGDATA_HOME}/om-server/OMS/workspace/conf/oms-config.ini | grep om_gateway
```

**Step 4** Run the **ping Management plane gateway** command on the active and standby management nodes and check whether the nodes are connected to the management plane gateway.

- If yes, go to [Step 5](#).
- If no, contact the network administrator to rectify the network fault. After the fault is rectified, go to [Step 5](#).

**Step 5** Run the following command on the active management node to obtain the domain name of the NTP server in the current environment:

This section uses **ntp.myhuaweicloud.com** as an example.

```
cat /opt/Bigdata_func/cloudinit/cloudinit_params | grep ntpserver
```

**Step 6** On the active management node, check the time difference between the new NTP server and the cluster. The unit is second.

For example, to check the time different with the NTP server at **ntp.myhuaweicloud.com**, run the **ntpdate -d ntp.myhuaweicloud.com** command. The following information is displayed:

```
6 Dec 15:16:10 ntpdate[2861453]: step time server 10.79.3.251 offset +2.118107 sec
```

In the preceding information, **+2.118107 sec** indicates the time offset. A positive value indicates that the NTP server time is earlier than the current cluster time. A negative value indicates the opposite.

•  **NOTE**

You can run the **ntpq -v** or **ntpq --version** command to query the NTP version. The command output may vary with the actual service environment.

– Output of the **ntpq -v** command:

```
10.1.1.112: ~# ntpq -v
ntpq - standard NTP query program - Ver. 4.2.4p8
```

– Output of the **ntpq --version** command:

```
10.1.1.112: ~# ntpq --version
ntpq 4.2.8p10@1.3728-o Mon Jun 6 08:01:59 UTC 2016 (1)
```

**Step 7** Check whether the absolute value of the time difference exceeds **150**.

- If yes, go to [Step 8](#).
- If no, perform [Step 10](#) as user **omm**.

**Step 8** Check whether the cluster can be stopped.

- If yes, stop upper-layer services and the cluster, and go to [Step 9](#).
- If no, no further action is required.

**Step 9** Check whether the time of the NTP server is slower than the time of the cluster.

- If yes, wait a period of the time difference obtained in [Step 6](#) after message **Operation successful** is displayed on the UI, perform [Step 11](#) as user **omm**.
- If no, after message **Operation successful** is displayed on the UI, perform [Step 11](#) as user **omm**.

**Step 10** Run the following command on the active management node to replace the NTP server:

```
sh ${BIGDATA_HOME}/om-server/om/bin/tools/modifyntp.sh --ntp_server_ip ntp.myhuaweicloud.com
```

 **NOTE**

The IP address of the NTP server cannot be set to the IP address of a node in the cluster. Otherwise, the service network between the node and the active/standby OMS node may be disconnected.

**Step 11** Run the following command on the active management node to forcibly synchronize time from the NTP server at **ntp.myhuaweicloud.com** immediately and replace the NTP server:

```
sh ${BIGDATA_HOME}/om-server/om/bin/tools/modifyntp.sh --ntp_server_ip ntp.myhuaweicloud.com --force_sync_time
```

 **NOTE**

- If the cluster is stopped, start the cluster after the NTP server is replaced.
- After the command for forcible time synchronization is executed, it takes about five minutes for time synchronization on cluster nodes.

----End

## 6.4.11 Modifying the OMS Service Configuration

Based on the security requirements of the user environment, you can modify the Kerberos and LDAP configurations in the OMS on FusionInsight Manager.

 **NOTE**

This section applies only to MRS 3.x or later.

### Impact on the System

After the OMS service configuration parameters are modified, the corresponding OMS module needs to be restarted. In this case, FusionInsight Manager cannot be used.

### Modifying the OMS Service Configuration

#### Modifying the okerberos configuration

**Step 1** Log in to FusionInsight Manager and choose **System > OMS**.

**Step 2** Locate the row that contains okerberos and click **Modify Configuration**.

**Step 3** Modify the parameters according to [Table 6-12](#).

**Table 6-12** okerberos parameters

| Parameter                   | Description                                                                                                                                                                                                                                                                                                            |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| KDC Timeout (ms)            | Timeout duration for an application to connect to Kerberos, in milliseconds. The value must be an integer.                                                                                                                                                                                                             |
| Max Retries                 | Maximum number of retries for an application to connect to Kerberos, in seconds. The value must be an integer.                                                                                                                                                                                                         |
| LDAP Timeout (ms)           | Timeout duration for Kerberos to connect to LDAP, in milliseconds.                                                                                                                                                                                                                                                     |
| LDAP Search Timeout (ms)    | Timeout duration for Kerberos to query user information in LDAP, in milliseconds.                                                                                                                                                                                                                                      |
| Kadmin Listening Port       | Port number of the Kadmin service.                                                                                                                                                                                                                                                                                     |
| KDC Listening Port          | Port number of the kinit service.                                                                                                                                                                                                                                                                                      |
| Kpasswd Listening Port      | Port number of the Kpasswd service.                                                                                                                                                                                                                                                                                    |
| Reset LDAP Account Password | Machine-machine users ( <b>cn=krbadmin,ou=Users,dc=hadoop,dc=com</b> and <b>cn=krbkdc,ou=Users,dc=hadoop,dc=com</b> ) used by Kerberos to access LDAP.<br>If this parameter is selected, the passwords will be replaced by random passwords.<br><b>NOTE</b><br>This parameter is available only in MRS 3.1.2 or later. |

**Step 4** Click **OK**.

In the displayed dialog box, enter the password of the current login user and click **OK**. In the displayed confirmation dialog box, click **OK**.

#### Modifying the oldap configuration

**Step 5** Locate the row that contains the oldap and click **Modify Configuration**.

**Step 6** Modify the parameters according to [Table 6-13](#).

**Table 6-13** OLDAP parameters

| Parameter                   | Description                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LDAP Listening Port         | Port number of the LDAP service.                                                                                                                                                                                                                                                                                                                    |
| Reset LDAP Account Password | Machine-machine users ( <b>cn=root,dc=hadoop,dc=com</b> and <b>cn=pg_search_dn,ou=Users,dc=hadoop,dc=com</b> ) used by LDAP for data management, synchronization, and status check.<br>If this parameter is selected, the passwords will be replaced by random passwords.<br><b>NOTE</b><br>This parameter is available only in MRS 3.1.2 or later. |

**Step 7** Click **OK**.

In the displayed dialog box, enter the password of the current login user and click **OK**. In the displayed confirmation dialog box, click **OK**.

 **NOTE**

To reset the password of the LDAP account, you need to restart ACS. The procedure is as follows:

1. Log in to the active management node as user **omm** using PuTTY, and run the following command to update the domain configuration:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/restart-RealmConfig.sh
```

The command is run successfully if the following information is displayed:

```
Modify realm successfully. Use the new password to log into FusionInsight again.
```

2. Run the **sh \$CONTROLLER\_HOME/sbin/acs\_cmd.sh stop** command to stop ACS.
3. Run the **sh \$CONTROLLER\_HOME/sbin/acs\_cmd.sh start** command to start ACS.

**Restarting the cluster****Step 8** Log in to FusionInsight Manager and restart the cluster by referring to [Restarting an MRS Cluster](#).

----End

## 6.4.12 Modifying MRS Manager Routing Table

When FusionInsight Manager is installed, two pieces of routing information are automatically created on the active management node. You can run the **ip rule list** command to view the routing information.

```
0:from all lookup local
32764:from all to 10.10.100.100 lookup ntp_rt #NTP routing information created by FusionInsight
Manager (this information is unavailable if no external NTP clock source is configured).
32765:from 192.168.0.117 lookup om_rt #OM routing information created by the FusionInsight Manager.
32766:from all lookup main
32767:from all lookup default
```

 NOTE

- If no external NTP server has been configured, only OM routing information **om\_rt** will be created.
- This section applies only to MRS 3.x or later.

If the routing information created by FusionInsight Manager conflicts with the routing information configured in the enterprise network planning, the cluster administrator can use **autoroute.sh** to disable or enable the routing information created by FusionInsight Manager.

## Impact on the System

After the routing information created by FusionInsight Manager is disabled and before the new routing information is set, FusionInsight Manager cannot be accessed but the clusters are running properly.

## Prerequisites

You have obtained the information about the route to be created.

## Disable the Routing Information Created by the System

- Step 1** Log in to the active management node of the cluster as the **omm** user. Run the following commands to disable the routing information created by the system:

```
cd ${BIGDATA_HOME}/om-server/om/sbin
./autoroute.sh disable
```

```
Deactivating Route.
Route operation (disable) successful.
```

- Step 2** View the execution result.

```
ip rule list
```

```
0:from all lookup local
32766:from all lookup main
32767:from all lookup default
```

- Step 3** Run the following command and enter the password of user **root** to switch to user **root**:

```
su - root
```

- Step 4** Run the following commands to manually create the routing information about the WS floating IP address:

```
ip route add Network segment of the WS floating IP address/Subnet mask of the WS floating IP address scope link src WS floating IP address dev NIC of the WS floating IP address table om_rt
```

```
ip route add default via Gateway of the WS floating IP address dev NIC of the WS floating IP address table om_rt
```

```
ip rule add from WS floating IP address table om_rt
```

Example:



```
ip route add 192.168.0.0/255.255.255.0 scope link src 192.168.0.117 dev eth0:ws table om_rt
```

```
ip route add default via 192.168.0.254 dev eth0:ws table om_rt
```

```
ip rule add from 192.168.0.117 table om_rt
```

#### NOTE

If IPv6 addresses are used, run the `ip -6 route add` command.

- Step 5** Run the following commands to manually create the NTP service routing information. Skip this step when no external NTP clock source is configured.

```
ip route add default via IP gateway of the NTP service dev NIC of the local IP address table ntp_rt
```

```
ip rule add to ntpIP table ntp_rt
```

*NIC of the local IP address* indicates the NIC that can communicate with the network segment where the NTP server is located.

Example:

```
ip route add default via 10.10.100.254 dev eth0 table ntp_rt
```

```
ip rule add to 10.10.100.100 table ntp_rt
```

- Step 6** View the execution result.

In the following example, if the command output contains `om_rt` and `ntp_rt`, the operation is successful.

#### ip rule list

```
0:from all lookup local
32764:from all to 10.10.100.100 lookup ntp_rt #This information is not displayed if no external NTP clock source is configured.
32765:from 192.168.0.117 lookup om_rt
32766:from all lookup main
32767:from all lookup default
```

----End

## Enable the Routing Information Created by the System

- Step 1** Log in to the active management node as user `omm`.

- Step 2** Run the following commands to enable the routing information created by the system:

```
cd ${BIGDATA_HOME}/om-server/om/sbin
```

```
./autoroute.sh enable
```

```
Activating Route.
Route operation (enable) successful.
```

- Step 3** View the execution result.

In the following example, if the command output contains `om_rt` and `ntp_rt`, the operation is successful.

### ip rule list

```
0:from all lookup local
32764:from all to 10.10.100.100 lookup ntp_rt #This information is not displayed if no external NTP clock
source is configured.
32765:from 192.168.0.117 lookup om_rt
32766:from all lookup main
32767:from all lookup default
```

----End

## 6.5 Managing MRS Cluster Components

### 6.5.1 Checking the Running Status of an MRS Cluster Component

After creating an MRS cluster, you can check the status of each service component and its role instance on the console or Manager. This will assist you in detecting any performance issues with the component.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

#### Viewing Component Status on the Console

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** On the MRS cluster details page, click **Components** to view the service operation status, health status, and configuration status.

**Figure 6-18** Checking component status

| Name      | Version | Operating Status | Health Status | Configuration Status | Roles                                                   | Operation                 |
|-----------|---------|------------------|---------------|----------------------|---------------------------------------------------------|---------------------------|
| ▼ Hadoop  | 3.3.1   |                  |               |                      |                                                         |                           |
| Spark     | 3.3.1   | Started          | Good          | Synchronized         | JobHistory 2 JDBCServer 2 SparkResource 3 IndexServer 2 | Start Stop Restart Delete |
| Flink     | 1.17.1  | Started          | Good          | Synchronized         | FlinkResource 2 FlinkServer 2                           | Start Stop Restart Delete |
| ZooKeeper | 3.6.1   | Started          | Good          | Synchronized         | quorumpeer 3                                            | Start Stop Restart Delete |

**Table 6-14** Component status description


| Status Type      | Status  | Description             |
|------------------|---------|-------------------------|
| Operating Status | Started | The service is started. |
|                  | Stopped | The service is stopped. |

| Status Type          | Status                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                      | Failed to start       | Failed to start the role instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|                      | Failed to stop        | Failed to stop the role instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|                      | Unknown               | Initial service status after the background system restarts.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Health Status        | Good                  | Indicates that all role instances in the service are running properly.                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                      | Faulty                | The running status of at least one role instance is <b>Faulty</b> or the status of the service on which the current service depends is abnormal.<br><br>If the running status of a service is <b>Faulty</b> , an alarm is generated. Rectify the fault based on the alarm information.                                                                                                                                                                                                            |
|                      | Unknown               | All role instances in the service are in the <b>Unknown</b> state.                                                                                                                                                                                                                                                                                                                                                                                                                                |
|                      | Restoring             | The background system is restarting the service.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                      | Partially Healthy     | The status of the service on which the service depends is abnormal, and APIs related to the abnormal service cannot be called by external systems.<br><br>HBase, Hive, Spark, and Loader may be in the <b>Partially Healthy</b> state. <ul style="list-style-type: none"> <li>If YARN is installed but is abnormal, HBase is in the <b>Partially Healthy</b> state.</li> <li>If HBase is installed but is abnormal, Hive, Spark, and Loader are in the <b>Partially Healthy</b> state.</li> </ul> |
| Configuration Status | Synchronized          | Indicates that the latest configuration takes effect.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|                      | Configuration expired | The latest configuration does not take effect after the parameter modification. You need to restart related services.                                                                                                                                                                                                                                                                                                                                                                             |
|                      | Configuration failed  | If a communication or read/write exception occurs during parameter configuration, you can use the synchronization configuration function to rectify the fault.                                                                                                                                                                                                                                                                                                                                    |
|                      | Configuring           | Parameters are being configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Status Type | Status  | Description                                      |
|-------------|---------|--------------------------------------------------|
|             | Unknown | Current configuration status cannot be obtained. |

**Step 4** Click a component name to go to the component details page and view the detailed running information about the component.

**Figure 6-19** Viewing cluster component details



**Step 5** Click **Instances** to view the detailed running information about each role instance in the service.

- The list of role instances shows all the instances for each role in the cluster, including their running and configuration status, hosts, and IP addresses.
- You can click an instance name to go to the instance details page and view the basic information, configuration file, instance logs, and monitoring metric graphs of the instance.

**Figure 6-20** Checking the status of cluster component instances

| Role     | Host Name            | OM IP Address | Business IP Address | Rack | Running status | Configuration Status |
|----------|----------------------|---------------|---------------------|------|----------------|----------------------|
| DataNode | node-group-10h4r0001 |               |                     |      | Good           | Synchronized         |
| DataNode | node-group-10h4r0002 |               |                     |      | Good           | Synchronized         |
| DataNode | node-group-1S4ha     |               |                     |      | Good           | Synchronized         |
| DataNode | node-group-10h4r0003 |               |                     |      | Good           | Synchronized         |
| HDFS     | node-master3a7q      |               |                     |      | Good           | Synchronized         |
| HDFS     | node-master1Bvy      |               |                     |      | Good           | Synchronized         |

**Table 6-15** Instance status

| Status Type    | Status         | Description                                            |
|----------------|----------------|--------------------------------------------------------|
| Running Status | Good           | The instance is running properly.                      |
|                | Bad            | The instance cannot run properly.                      |
|                | Decommissioned | The instance is out of service.                        |
|                | Not started    | The instance is stopped.                               |
|                | Unknown        | The initial status of the instance cannot be detected. |
|                | Starting       | The instance is being started.                         |
|                | Stopping       | The instance is being stopped.                         |

| Status Type          | Status                | Description                                                                                                                                                    |
|----------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                      | Restoring             | An exception may occur in the instance and the instance is being automatically rectified.                                                                      |
|                      | Decommissioning       | The instance is being decommissioned.                                                                                                                          |
|                      | Recommissioning       | The instance is being recommissioned.                                                                                                                          |
|                      | Failed to start       | The instance fails to be started.                                                                                                                              |
|                      | Failed to stop        | The instance fails to be stopped.                                                                                                                              |
| Configuration Status | Synchronized          | The latest configuration takes effect.                                                                                                                         |
|                      | Configuration expired | The latest configuration does not take effect after the parameter modification. You need to restart related services.                                          |
|                      | Configuration failed  | If a communication or read/write exception occurs during parameter configuration, you can use the synchronization configuration function to rectify the fault. |
|                      | Configuring           | Parameters are being configured.                                                                                                                               |
|                      | Unknown               | Current configuration status cannot be obtained.                                                                                                               |

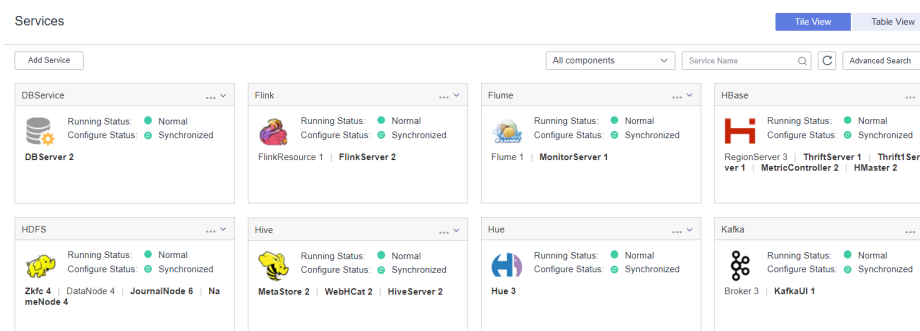
----End

## Checking the Component Status on Manager

**Step 1** Log in to Manager and choose **Cluster > Services** to access the component management page.

The service list displays the running status, configuration status, role type, and number of instances of each component.

**Figure 6-21** Checking the status of cluster components



 **NOTE**

On the Manager of MRS 2.x or earlier, click **Services** to open the component management page.


**Table 6-16** Manager component status description

| Status Type          | Status                                                                                                                                | Description                                                                                                                                                    |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Running Status       | Normal                                                                                                                                | The component is running properly.                                                                                                                             |
|                      | Faulty                                                                                                                                | The component cannot work properly.                                                                                                                            |
|                      | Partially Healthy                                                                                                                     | Some enhanced functions of the component cannot work properly.                                                                                                 |
|                      | Not started                                                                                                                           | The component is stopped.                                                                                                                                      |
|                      | Unknown                                                                                                                               | The initial status of the component cannot be detected.                                                                                                        |
|                      | Starting                                                                                                                              | The component is being started.                                                                                                                                |
|                      | Stopping                                                                                                                              | The component is being stopped.                                                                                                                                |
|                      | Failed to start                                                                                                                       | The component fails to be started.                                                                                                                             |
|                      | Failed to stop                                                                                                                        | The component fails to be stopped.                                                                                                                             |
| Configuration Status | Synchronized                                                                                                                          | The latest configuration takes effect.                                                                                                                         |
|                      | <ul style="list-style-type: none"> <li>Expired (Manager 2.x or earlier)</li> <li>Expired (Manager 3.x and later)</li> </ul>           | The latest configuration does not take effect after the parameter modification. You need to restart related services.                                          |
|                      | Failed                                                                                                                                | If a communication or read/write exception occurs during parameter configuration, you can use the synchronization configuration function to rectify the fault. |
|                      | <ul style="list-style-type: none"> <li>Configuring (Manager 2.x or earlier)</li> <li>Synchronizing (Manager 3.x and later)</li> </ul> | Parameters are being configured.                                                                                                                               |
|                      | Unknown                                                                                                                               | Current configuration status cannot be obtained.                                                                                                               |

**Step 2** Click a component name to view its details.

**Step 3** Click **Instances** to view the detailed running information about each role instance in the service.

**Figure 6-22** Checking the status of cluster component instances



- The list of role instances shows all the instances for each role in the cluster, including their running and configuration status, hosts, and IP addresses.
- You can click an instance name to go to the instance details page and view the basic information, configuration file, instance logs, and monitoring metric graphs of the instance.

**Table 6-17** Manager instance status description (3.x and later versions)

| Status Type          | Status                            | Description                                                                                                                                                    |
|----------------------|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Running Status       | Normal                            | The instance is running properly.                                                                                                                              |
|                      | Faulty                            | The instance cannot run properly.                                                                                                                              |
|                      | Decommissioned                    | The instance is out of service.                                                                                                                                |
|                      | Not started                       | The instance is stopped.                                                                                                                                       |
|                      | Unknown                           | The initial status of the instance cannot be detected.                                                                                                         |
|                      | Starting                          | The instance is being started.                                                                                                                                 |
|                      | Stopping                          | The instance is being stopped.                                                                                                                                 |
|                      | Restoring                         | An exception may occur in the instance and the instance is being automatically rectified.                                                                      |
|                      | Decommissioning                   | The instance is being decommissioned.                                                                                                                          |
|                      | Recommissioning                   | The instance is being recommissioned.                                                                                                                          |
|                      | Failed to start                   | The instance fails to be started.                                                                                                                              |
| Failed to stop       | The instance fails to be stopped. |                                                                                                                                                                |
| Configuration Status | Synchronized                      | The latest configuration takes effect.                                                                                                                         |
|                      | Expired                           | The latest configuration does not take effect after the parameter modification. You need to restart related services.                                          |
|                      | Failed                            | If a communication or read/write exception occurs during parameter configuration, you can use the synchronization configuration function to rectify the fault. |

| Status Type | Status        | Description                                      |
|-------------|---------------|--------------------------------------------------|
|             | Synchronizing | Parameters are being configured.                 |
|             | Unknown       | Current configuration status cannot be obtained. |

**Table 6-18** Manager instance status description (2.x and earlier versions)

| Status Type          | Status            | Description                                                                                                                                                       |
|----------------------|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operating Status     | Started           | The role instance has been started.                                                                                                                               |
|                      | Stopped           | The role instance has been stopped.                                                                                                                               |
|                      | Failed to start   | Failed to start the role instance.                                                                                                                                |
|                      | Failed to stop    | Failed to stop the role instance.                                                                                                                                 |
|                      | Decommissioning   | The role instance is being decommissioned.                                                                                                                        |
|                      | Decommissioned    | The role instance has been decommissioned.                                                                                                                        |
|                      | Recommissioning   | The role instance is being recommissioned.                                                                                                                        |
|                      | Unknown           | Initial role instance status after the background system restarts.                                                                                                |
| Health Status        | Good              | The role instance is running properly.                                                                                                                            |
|                      | Restoring         | The background system is restarting a role instance.                                                                                                              |
|                      | Bad               | The instance role is experiencing an abnormality, such as the inability to access a port due to a non-existent PID.                                               |
|                      | Unknown           | The host where a role instance resides does not connect to the background system.                                                                                 |
|                      | Partially Healthy | The role instance is partially running properly.                                                                                                                  |
| Configuration Status | Synchronized      | The latest configuration takes effect.                                                                                                                            |
|                      | Expired           | The latest configuration does not take effect after the parameter modification. Related services need to be restarted.                                            |
|                      | Failed            | The communication is incorrect or data cannot be read or written during the parameter configuration. Click <b>Synchronize Configuration</b> to rectify the fault. |



| Status Type | Status      | Description                                      |
|-------------|-------------|--------------------------------------------------|
|             | Configuring | Parameters are being configured.                 |
|             | Unknown     | Current configuration status cannot be obtained. |

----End

## 6.5.2 Starting and Stopping an MRS Cluster Component

You can stop a service component in the MRS cluster as required. After the component is stopped, it will no longer offer any services.

- Stop the services or abnormal services.
- Start the service in the **Stopped**, **Stop Failed**, or **Failed to Start** state to use the service again.
- Restart abnormal services or configure expired services to restore or enable the services.

### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Impact on the System

Services within the system rely on one another. Therefore, when you initiate, halt, or reboot a service, any services that depend on it will be impacted.

- If a service is to be started, the lower-layer services dependent on it must be started first.
- If a service is stopped, the upper-layer services dependent on it are unavailable.
- If a service is restarted, the running upper-layer services dependent on it must be restarted.

### Starting or Stopping a Component on the Console

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** On the MRS cluster details page, click **Components**.
- Step 4** Locate the row that contains the target service and click **Stop**.

Click **Start** to start the service if needed.

----End


## Starting or Stopping a Component on Manager

For MRS 3.x or later:

**Step 1** Log in to Manager and choose **Cluster > Services** to access the service management page.

The displayed service list contains all installed services in the cluster. If the tile view is selected, the services will be displayed in pane style. If you select the list view, the services will be displayed in a table.

In this section, the **Tile View** is used by default.

**Step 2** In the upper right corner of the specified service pane, choose  > **Stop Service**, enter the password, and click **OK**.

You can also click the name of the service to be operated, click **Stop Service** in the upper right corner of the **Dashboard** page, enter the password, confirm the operation impact, and click **OK**.

----End

For MRS 2.x or earlier:

**Step 1** Log in to Manager and click **Services**.

**Step 2** Locate the row that contains the target service and click **Stop**.

Click **Start** to start the service if needed.

----End

## 6.5.3 Restarting an MRS Cluster Component

To apply configuration changes to a big data component, you must restart it. However, using the common restart mode will restart all services or instances at once, which can cause service interruption.

To ensure that services are not affected during service restart, you can restart services or instances in batches by rolling restart. For instances in active/standby mode, a standby instance is restarted first and then an active instance is restarted.

A rolling restart takes a longer time and may affect service throughput and performance.

For details about whether services and instances in the current MRS cluster support rolling restart and the rolling restart parameters, see [Component Restart Reference Information](#).

### Restrictions

- Perform a rolling restart during off-peak hours.
  - If the service throughput of the Kafka service is high (over 100 MB/s) during a rolling restart, the restart will fail.

- To avoid RegionServer restart failures caused by heavy loads during an HBase rolling restart, increase the number of handles if the requests per second of each RegionServer on the native interface exceed 10,000.
- Before restarting, check the current number of requests in HBase. If the number of requests on the native interface for each RegionServer is over 10,000, increase the number of handles to prevent overloading.
- If the number of Core nodes in a cluster is less than six, services may be affected for a short period of time.
- Preferentially perform a rolling instance or service restart and select **Only restart instances whose configurations have expired**.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

## Restarting Cluster Components

**Step 1** Access the MRS cluster component management page.

- Log in to the MRS console and click the cluster name to go to the cluster details page. Click **Components**.
- If you are using the Manager of MRS 3.x and later versions, log in to Manager and choose **Cluster > Services**.
- If you are using the Manager of MRS 2.x and earlier versions, log in to Manager and click **Services**.

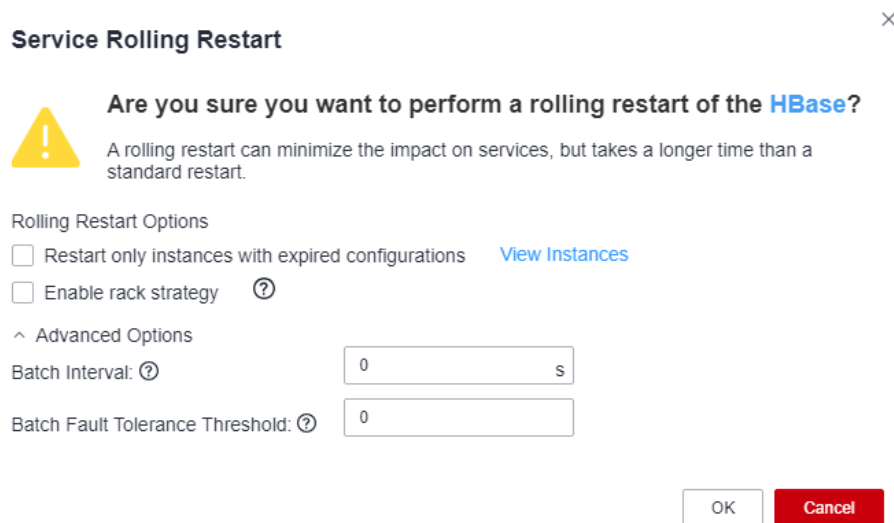
**Step 2** Click the name of the target component to go to the details page.

**Step 3** On the service details page, expand the **More** drop-down list and select **Restart Service** or **Service Rolling Restart**.

**Step 4** Enter the user password (required when you perform operations on Manager), confirm the operation impact, and click **OK** to restart the system.

If you select rolling restart, set parameters listed in [Table 6-19](#). (Required parameters may vary by version, set parameters based on the actual GUI.)

**Figure 6-23** Performing a rolling restart on Manager



**Table 6-19** Rolling restart configuration parameters

| Parameter                                          | Description                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Restart only instances with expired configurations | Whether to restart only the modified instances in a cluster. The name of this parameter may be different in other versions.                                                                                                                                                                                                                    |
| Enable rack strategy                               | Whether to enable the concurrent rack rolling restart strategy. This parameter takes effect only for roles that meet the rack rolling restart strategy. (The roles support rack awareness, and instances of the roles belong to two or more racks.)<br><br>This parameter can be set only when a rolling restart is performed on HDFS or YARN. |

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data Nodes to Be Batch Restarted | <p>Number of instances that are restarted in each batch when the batch rolling restart strategy is used. The default value is <b>1</b>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>This parameter is valid only when the batch rolling restart strategy is used and the instance type is DataNode.</li> <li>This parameter is invalid when the rack strategy is enabled. In this case, the cluster uses the maximum number of instances (<b>20</b> by default) configured in the rack strategy as the maximum number of instances that are concurrently restarted in a rack.</li> <li>This parameter can be set only when a rolling restart is performed on some components, such as HDFS, HBase, YARN, Kafka, Storm, and Flume. The actual value displayed on the GUI prevails.</li> <li>The number of concurrent RegionServer rolling restarts of HBase cannot be manually configured. It is automatically adjusted based on the number of RegionServer nodes. The adjustment rules are as follows: If the number of nodes is less than 30, one node will be added in each batch. For node counts less than 300, two nodes will be added in each batch. If the node count exceeds 300 (including 300 nodes), each batch will add 1% (rounded down) of the total nodes.</li> </ul> |
| Batch Interval                   | <p>Interval between two batches of instances to be roll-restarted. The default value is <b>0</b>.</p> <p>Setting the batch interval parameter can increase the stability of the big data component process during the rolling restart.</p> <p>You are advised to set this parameter to a non-default value, for example, <b>10</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Decommissioning Timeout Interval | <p>Decommissioning waiting time of a role instance during a rolling restart. This parameter can be set only when a rolling restart is performed on Hive or Spark.</p> <p>Some roles (such as HiveServer and JDBCServer) stop providing services before the rolling restart. Stopped instances cannot be connected to new clients. Existing connections will be completed after a period of time. An appropriate timeout interval can ensure service continuity.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Batch Fault Tolerance Threshold  | <p>Tolerance times when the rolling restart of instances fails to be batch executed. The default value is <b>0</b>, which indicates that the rolling restart task ends after any batch of instances fails to restart.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

----End

## Component Restart Reference Information

**Table 6-20** provides services and instances that support or do not support rolling restart in the MRS cluster.

**Table 6-20** Services and instances that support or do not support rolling restart

| Service    | Instance           | Rolling Restart |
|------------|--------------------|-----------------|
| Alluxio    | AlluxioJobMaster   | Yes             |
|            | AlluxioMaster      |                 |
| ClickHouse | ClickHouseServer   | Yes             |
|            | ClickHouseBalancer |                 |
| CDL        | CDLConnector       | Yes             |
|            | CDLService         |                 |
| Flink      | FlinkResource      | No              |
|            | FlinkServer        |                 |
| Flume      | Flume              | Yes             |
|            | MonitorServer      |                 |
| Guardian   | TokenServer        | Yes             |
| HBase      | HMaster            | Yes             |
|            | RegionServer       |                 |
|            | ThriftServer       |                 |
|            | RETServert         |                 |
| HetuEngine | HSBroker           | Yes             |
|            | HSConsole          |                 |
|            | HSFabric           |                 |
|            | QAS                |                 |
| HDFS       | NameNode           | Yes             |
|            | Zkfc               |                 |
|            | JournalNode        |                 |
|            | HttpFS             |                 |
|            | DataNode           |                 |
| Hive       | MetaStore          | Yes             |
|            | WebHCat            |                 |
|            | HiveServer         |                 |
| Hue        | Hue                | No              |
| Impala     | Impalad            | No              |

| Service   | Instance         | Rolling Restart |
|-----------|------------------|-----------------|
|           | StateStore       |                 |
|           | Catalog          |                 |
| IoTDB     | IoTDBServer      | Yes             |
| Kafka     | Broker           | Yes             |
|           | KafkaUI          | No              |
| Kudu      | KuduTserver      | Yes             |
|           | KuduMaster       |                 |
| Loader    | Sqoop            | No              |
| MapReduce | JobHistoryServer | Yes             |
| Oozie     | oozie            | No              |
| Presto    | Coordinator      | Yes             |
|           | Worker           |                 |
| Ranger    | RangerAdmin      | Yes             |
|           | UserSync         |                 |
|           | TagSync          |                 |
| Spark     | JobHistory       | Yes             |
|           | JDBCServer       |                 |
|           | SparkResource    |                 |
| Storm     | Nimbus           | Yes             |
|           | UI               |                 |
|           | Supervisor       |                 |
|           | Logviewer        |                 |
| Tez       | TezUI            | No              |
| Yarn      | ResourceManager  | Yes             |
|           | NodeManager      |                 |
| ZooKeeper | Quorumpeer       | Yes             |

**Table 6-21** lists the instance startup duration.

**Table 6-21** Restart duration for reference

| Service    | Restart Duration | Startup Duration                                                                                                                                       | Remarks                                                                                                                                                                                                                                         |
|------------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IoTDB      | 3min             | IoTDBServer: 3 min                                                                                                                                     | -                                                                                                                                                                                                                                               |
| CDL        | 2min             | <ul style="list-style-type: none"> <li>CDLConnector: 1 min</li> <li>CDLService: 1 min</li> </ul>                                                       | -                                                                                                                                                                                                                                               |
| ClickHouse | 4min             | <ul style="list-style-type: none"> <li>ClickHouseServer: 2 min</li> <li>ClickHouseBalancer: 2 min</li> </ul>                                           | -                                                                                                                                                                                                                                               |
| HDFS       | 10min+x          | <ul style="list-style-type: none"> <li>NameNode: 4 min + x</li> <li>DataNode: 2 min</li> <li>JournalNode: 2 min</li> <li>Zkfc: 2 min</li> </ul>        | <p>x indicates the NameNode metadata loading duration. It takes about 2 minutes to load 10,000,000 files. For example, x is 10 minutes for 50 million files.</p> <p>The startup duration fluctuates with reporting of DataNode data blocks.</p> |
| Yarn       | 5min+x           | <ul style="list-style-type: none"> <li>ResourceManager: 3 min + x</li> <li>NodeManager: 2 min</li> </ul>                                               | x indicates the time required for restoring ResourceManager reserved tasks. It takes about 1 minute to restore 10,000 reserved tasks.                                                                                                           |
| MapReduce  | 2min+x           | JobHistoryServer: 2 min + x                                                                                                                            | x indicates the scanning duration of historical tasks. It takes about 2.5 minutes to scan 100,000 tasks.                                                                                                                                        |
| ZooKeeper  | 2min+x           | quorumpeer: 2 min + x                                                                                                                                  | x indicates the duration for loading znodes. It takes about 1 minute to load 1 million znodes.                                                                                                                                                  |
| Hive       | 3.5min           | <ul style="list-style-type: none"> <li>HiveServer: 3 min</li> <li>MetaStore: 1 min 30s</li> <li>WebHcat: 1 min</li> <li>Hive service: 3 min</li> </ul> | -                                                                                                                                                                                                                                               |



| Service | Restart Duration | Startup Duration                                                                                                                           | Remarks                                                                                                               |
|---------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Spark2x | 5min             | <ul style="list-style-type: none"><li>● JobHistory2x: 5 min</li><li>● SparkResource 2x: 5 min</li><li>● JDBCServer2x: 5 min</li></ul>      | -                                                                                                                     |
| Flink   | 4min             | <ul style="list-style-type: none"><li>● FlinkResource: 1 min</li><li>● FlinkServer: 3 min</li></ul>                                        | -                                                                                                                     |
| Kafka   | 2min+x           | <ul style="list-style-type: none"><li>● Broker: 1 min + x</li><li>● Kafka UI: 5 min</li></ul>                                              | x indicates the data restoration duration. It takes about 2 minutes to start 20,000 partitions for a single instance. |
| Storm   | 6min             | <ul style="list-style-type: none"><li>● Nimbus: 3 min</li><li>● UI: 1 min</li><li>● Supervisor: 1 min</li><li>● Logviewer: 1 min</li></ul> | -                                                                                                                     |
| Flume   | 3min             | <ul style="list-style-type: none"><li>● Flume: 2 min</li><li>● MonitorServer: 1 min</li></ul>                                              | -                                                                                                                     |
| Doris   | 2 min            | <ul style="list-style-type: none"><li>● FE: 1min</li><li>● BE: 1min</li><li>● DBroker: 1min</li></ul>                                      | -                                                                                                                     |

## 6.5.4 Adding and Deleting an MRS Cluster Component

When creating an MRS cluster, you can select service components to be included in the cluster. After the cluster is created, administrators can manually add or delete components for MRS on the console.

### NOTE

MRS 3.1.2-LTS.3 and later versions allow you to add and delete components.

### Prerequisites

- You have configured permissions for the user group to which the IAM users belong.

Adding or deleting a service in a cluster is a high-risk operation. Bind the MRS FullAccess, MRS Administrator, Server Administrator, Tenant Guest, MRS

Administrator, or Tenant Administrator policy to the user group before you perform this operation.

For details about the permissions, see [Synchronizing IAM Users to MRS](#).

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.

## Adding or Deleting a Cluster Component

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the cluster details page, choose **Components** and click **Add Service**.

**Step 4** In the service list, select the services to be added and click **Next**.

### NOTE

- When you add a service, the underlying services on which the service depends are automatically selected. You can add multiple services at the same time.
- You can add services only on nodes or node groups that are in normal state.
- Components (MapReduce, YARN, and HDFS) in the Hadoop service cannot be added separately.
- If Hadoop is not installed in a cluster, after Hadoop is added, you need to refresh the console and synchronize IAM users again to submit jobs on the job management page.
- After the Spark2x/Spark component is added, if you need to perform operations on Spark SQL on the Hue web UI, restart the Hue service first.

**Step 5** On the **Topology Adjustment** page, select the nodes where the service is to be deployed. For details about the deployment scheme, see [Table 2-5](#).

**Step 6** Confirm the operation impact and click **OK**. After the service is added, you can view the added service on the **Components** page.

**Step 7** Locate the row that contains the service and click **Delete** to delete a service.

Enter **DELETE** in the displayed **Delete Service** dialog box and click **OK** to confirm the deletion.

### NOTE

- If the service to be deleted has upper-layer dependencies, the service cannot be deleted. Only one service can be deleted at a time.
- You can delete installed services except Hadoop (HDFS, YARN, and MapReduce), Ranger, DBService, KrbServer, LdapServer, and meta services.
- Before deleting a service, back up the service data to prevent data loss.

----End

## 6.5.5 Modifying the Configuration Parameters of an MRS Cluster Component

To meet actual service requirements, quickly view and modify default service configurations in MRS.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

## Impact on the System

- After configuring the HBase, HDFS, Hive, Spark, YARN and MapReduce service attributes, you need to download and update the client configuration file again.
- The parameters of DBService cannot be modified when only one DBService role instance exists in the cluster.
- After configuring properties of a service, restart the service if the service status is **Expired**. The service is unavailable during the restart.
- After the service configuration parameters are modified and then take effect after restart, you need to download and install the client again or download the configuration file to update the client.

## Modifying Parameters on the Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS cluster details page, click **Components**.

Figure 6-24 Components tab page

| Name      | Version    | Operating Status | Health Status |
|-----------|------------|------------------|---------------|
| Hadoop    | 3.1.1      |                  |               |
| Spark2x   | 2.4.5      | Started          | Faulty        |
| HBase     | 2.2.3      | Started          | Good          |
| Hive      | 3.1.0      | Started          | Faulty        |
| Hue       | 4.7.0      | Started          | Good          |
| Kafka     | 2.11-2.4.0 | Started          | Good          |
| Flume     | 1.9.0      | Started          | Good          |
| Flink     | 1.12.0     | Started          | Good          |
| Oozie     | 5.1.0      | Started          | Faulty        |
| ZooKeeper | 3.5.6      | Started          | Good          |

**Step 4** Select the target service from the service list.


**Step 5** Click **Service Configuration**. Switch **Basic** to **All**. All configuration parameters of the service are displayed in the navigation tree. The service name and role names are displayed from upper to lower in the navigation tree.

**NOTE**

The parameters under the service node are service-level configuration parameters, and the parameters under the role node are role-level configuration parameters. For details, see [Modifying MRS Role Instance Parameters](#).

**Step 6** In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

Figure 6-25 Modifying component configuration parameters



If you want to cancel the modification of a parameter value, click to restore it.

**NOTE**

- Select a port parameter value from the value range on the right. Ensure that all parameter values in the same service are within the value range and are unique. Otherwise, the service fails to be started.
- Configure parameters based on the information provided in the configuration description.

**Step 7** Click **Save Configuration**, save the parameters as prompted, and restart the service.

**NOTE**

In versions earlier than MRS 3.x, to update the queue configuration of YARN without restarting the service, choose **More > Refresh Queue** on the **Service Status** tab page.

----End

## Modifying Component Parameters on Manager

For clusters of MRS 3.x or later:

**Step 1** Log in to Manager and choose **Cluster > Services**.

**Step 2** Click the name of the service to be operated and click **Configurations**.

The **Basic Configurations** page is displayed by default. To modify more parameters, click **All Configurations**. The navigation tree displays all configuration parameters of the service. The level-1 nodes in the navigation tree are service names or role names. The parameter category is displayed after the level-1 node is expanded.

As shown in the following figure, **LdapServer** indicates the service name, indicating that the configuration applies to the entire service; **SlapdServer** indicates the role name, indicating that the configuration applies to all instances in the role.

**Figure 6-26** Configuration parameter navigation tree

**NOTE**

The parameters under the service node are service-level configuration parameters, and the parameters under the role node are role-level configuration parameters. For details, see [Modifying MRS Role Instance Parameters](#).

**Step 3** In the navigation tree, select the specified parameter category and change the parameter values on the right.

If you are not sure about the location of a parameter, you can enter the parameter name in search box in the upper right corner. The Manager searches for the parameter in real time and displays the result.



 **NOTE**

- Select a port parameter value from the value range on the right. Ensure that all parameter values in the same service are within the value range and are unique. Otherwise, the service fails to be started.
- Configure parameters based on the information provided in the configuration description.

**Step 4** Click **Save**. In the confirmation dialog box, click **OK**.

Wait until the message "Operation succeeded" is displayed. Click **Finish**. The configuration is modified.

 **NOTE**

- To update the queue configuration of the YARN service without restarting service, choose **More > Refresh Queue** to update the queue for the configuration to take effect.
- You can both download and upload configuration files. To configure Flume parameters, edit the **flume.config.file** file and upload it to apply the changes. After a configuration file is uploaded, the old file will be overwritten. If the configuration is not saved and the service is restarted, the configuration does not take effect. Save the configuration in time.
- If you need to restart the service for the configuration to take effect after modifying service configuration parameters, choose **More > Restart Service** in the upper right corner of the service page.
- If  is displayed before a parameter, this parameter takes effect dynamically. After the configuration is saved, the parameter value is automatically updated to the configuration file.  is supported only in MRS 3.2.0 or later.

----End

For clusters of MRS 2.x or earlier:

**Step 1** Log in to Manager and click **Services**.

**Step 2** Select the target service from the service list.

**Step 3** Click **Service Configuration**.

**Step 4** Set **Type** to **All**. All configuration parameters of the service are displayed in the navigation tree. The root nodes from top down in the navigation tree represent the service names and role names.

**Step 5** In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

To cancel the change to a parameter value, click .

**Step 6** Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK** to restart the service.

After the system displays "Operation succeeded", click **Finish**. The service is started.

 NOTE

To update the queue configuration of the YARN service without restarting service, choose **More > Refresh Queue** to update the queue for the configuration to take effect.

----End

## 6.5.6 Viewing the Modified Component Configuration Parameters of an MRS Cluster

MRS allows you to view the changes of service configuration parameters in a cluster with one click, helping you quickly locate faults and improve configuration management efficiency.

On the Manager page of the MRS 3.x cluster, you can easily see the non-initial default values for each service, non-unified values for the same role instance, historical records of cluster configuration changes, and expired configuration parameters in the cluster.

In MRS 3.x, O&M personnel can view environment variables and role configurations on Manager to quickly verify the accuracy of configuration items or access hidden configuration items.



### Viewing the Modified Component Configuration Parameters

**Step 1** Log in to the Manager page of the cluster.

**Step 2** Choose **Cluster > Configurations**.


**Step 3** Select an operation page based on the scenario.


- View all non-default values.
  - a. Click **All Non-default Values**. The system displays the parameters whose values are different from the default values configured for each service, role, or instance in the current cluster.

You can click  next to a parameter value to quickly restore the value to the default one. You can click  to view the historical modification records of the parameter.

If there are a large number of parameters to configure, you can filter the parameters in the filter box in the upper right corner of the page or enter keywords in the search box.

- b. To change the values of the parameters, change the values according to the parameter description and click **Save**. In the dialog box that is displayed, click **OK**.
- View all non-uniform values.
    - a. Click **All Non-uniform Values**. The system displays parameters with different role, service, instance group, or instance configurations in the current cluster.

You can click  next to a parameter value and view the differences in the dialog box that is displayed.

- b. To change the value of a parameter, click  to cancel the configuration difference or manually adjust the parameter value, click **OK**, and then click **Save**. In the dialog box that is displayed, click **OK**.
- View the expired configuration items.
  - a. Click **Expired Configurations**. Expired configuration items in the current cluster are displayed.
  - b. You can filter services using the service filter box in the upper part of the page to view expired configurations of different services. Alternatively, you can enter keywords in the search box.
  - c. Expired configuration items do not take effect completely. Restart the services or instances whose configurations have expired in a timely manner.
- View historical configuration records.
  - a. Click **Historical Configurations**. The historical configuration change records of the current cluster are displayed. You can view details about parameter value changes, including the service to which the parameter belongs, parameter values before and after the modification, and parameter files.
  - b. To restore a configuration change, click **Restore Configuration** in the **Operation** column of the target record. In the dialog box that is displayed, click **OK**.

If **the same configuration of the same service instance** is modified for multiple times, the configuration can only be restored to the latest modification.

 **NOTE**


Some configuration items take effect only after the corresponding services are restarted. After the configurations are saved, restart the services or instances whose configurations have expired in a timely manner.

----End

## Viewing the Role Instance Configuration File

- Step 1** Log in to the cluster Manager.
- Step 2** Choose **Cluster > Services**.
- Step 3** On the page that is displayed, click **Instances**.
- Step 4** Click the name of the target instance. In the **Configuration File** area on the **Instance Status** page, the configuration file list of the instance is displayed.



**Figure 6-27** Viewing the instance configuration file

**Step 5** Click the name of the configuration file to be viewed to view the parameter values in the configuration file.

To obtain the configuration file, download the configuration file to the local PC.

 **NOTE**

If a node in the cluster is faulty, the configuration file cannot be viewed. Rectify the fault before viewing the configuration file again.

----End

## 6.5.7 Synchronizing MRS Component Configuration Parameters

If you find that the **status of some components or instances expires or fails**, you can synchronize the configuration to restore the configuration state. If all services in the cluster are in the **Failed** state, synchronize the cluster configuration with the background configuration.

- If all services in the cluster are in the **Configuration failed** state, synchronize the cluster configuration with the background configuration.
- If all services in the cluster are in the **Configuration failed** state, synchronize the service configuration with the background configuration.

### Impact on the System

- After synchronizing cluster or service configurations, restart the services whose configurations have expired. These services are unavailable during the restart.
- After synchronizing the role instance configuration, restart the role instance whose configuration has expired. The role instance is unavailable during the restart.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

## Synchronizing Configurations on the Console


### Synchronizing cluster configurations

**Step 1** Log in to the MRS console and click the name of the target cluster.

**Step 2** On the cluster details page, choose **Configuration > Synchronize Configuration** in the upper right corner.

This operation applies only to MRS 2.x and earlier versions.

**Figure 6-28** Synchronizing configurations (using MRS 1.9.2 as an example)



**Step 3** In the displayed dialog box, select "Restart services and instances whose configuration have expired" and click **OK** to restart the service whose configurations have expired.

When **Operation successful** is displayed, click **Finish**. The cluster is started.

----End

### Synchronizing component configurations

**Step 1** Log in to the MRS console and click the name of the target cluster.

**Step 2** On the MRS cluster details page, click **Components**.

**Step 3** Select the target service from the service list.

**Step 4** On the **Service Status** page, choose **More > Synchronize Configuration** and operate as prompted.

----End

### Synchronizing role instance configurations

**Step 1** Log in to the MRS console and click the name of the target cluster.


**Step 2** On the MRS cluster details page, click **Components**.

**Step 3** Select the target service from the service list.

**Step 4** On the service details page, click **Instances**.

**Step 5** Click the target role instance from the role instance list.

**Step 6** In the upper part of the role instance status and metric information, choose **More > Synchronize Configuration** and synchronize the configuration as prompted.

**Figure 6-29** Configuration synchronization

----End

## Synchronizing Configurations Using Manager

### Synchronizing cluster configurations

**Step 1** Log in to Manager.

- MRS 3.x: Choose **Cluster > Dashboard**.  
For MRS 3.3.0 or later, the **Cluster > Dashboard** page has been removed from Manager. You can choose **More** in the upper right corner of the **Homepage** to access cluster maintenance and management functions.
- For MRS 2.x and earlier versions, click **Services**.

**Step 2** On this page, choose **More > Synchronize Configuration**.

**Step 3** Perform operations based on the cluster version.

- MRS 2.x and earlier versions: Perform identity authentication, select "Restart services and instances whose configuration have expired", and click **OK**. No further action is required.
- MRS 3.x and later versions: Click **OK**. After the synchronization is complete, perform the operations in [Step 4](#) to restart the instance whose configurations have expired.

**Step 4** Restart the instance whose configurations have expired. (Perform this operation only for MRS 3.x and later versions.)

1. On Manager, choose **Cluster > Dashboard**.
2. Choose **More > Restart Configuration-Expired Instances**.
3. In the dialog box that is displayed, enter the password of the current login user and click **OK**.
4. In the displayed dialog box, click **OK**.

Click **View Instance** to open the list of all expired instances and confirm that the instances have been restarted.

----End

### Synchronizing component configurations

**Step 1** Log in to Manager.

- MRS 3.x: Choose **Cluster > Services**.
- For MRS 2.x and earlier versions, click **Services**.

**Step 2** Click the specified service name and choose **More > Synchronize Configuration**.

**Step 3** Perform operations based on the cluster version.

- MRS 2.x and earlier versions: Perform identity authentication, select "Restart services and instances whose configuration have expired", and click **OK**. No further action is required.
- MRS 3.x and later versions: Click **OK**. After the synchronization is complete, perform the operations in [Step 4](#) to restart the instance whose configurations have expired.

**Step 4** Restart the instance whose configurations have expired. (Perform this operation only for MRS 3.x and later versions.)

1. On Manager, choose **Cluster > Dashboard**.
2. Choose **More > Restart Configuration-Expired Instances**.
3. In the dialog box that is displayed, enter the password of the current login user and click **OK**.
4. In the displayed dialog box, click **OK**.

Click **View Instance** to open the list of all expired instances and confirm that the instances have been restarted.

----End

### Synchronizing role instance configurations

**Step 1** Log in to Manager.

- MRS 3.x: Choose **Cluster > Services**.
- For MRS 2.x and earlier versions, click **Services**.

**Step 2** Click the specified service name and click **Instances**.

**Step 3** Click the target role instance from the role instance list.

**Step 4** On the role instance information page, choose **More > Synchronize Configuration**.

**Step 5** Perform operations based on the cluster version.

- MRS 2.x and earlier versions: In the displayed dialog box, select "Restart services and instances whose configuration have expired" and click **OK**.
- MRS 3.x and later versions: Click **OK**. After the synchronization is complete, perform the operations in [Step 4](#) to restart the instance whose configurations have expired.

**Step 6** Restart the instance whose configurations have expired. (Perform this operation only for MRS 3.x and later versions.)

1. Click **Instances**.
2. Select the instance to be restarted and choose **More > Restart Instance**.
3. In the dialog box that is displayed, enter the password of the current login user and click **OK**.
4. In the displayed dialog box, confirm the impact and click **OK**.

Click **View Instance** to open the list of all expired instances and confirm that the instances have been restarted.

----End

## 6.5.8 Adding Custom MRS Component Parameters

The big data components in the cluster fully support all parameters from the open source community. You can modify configuration parameters on the MRS console or MRS Manager for common scenarios, but some client components may not support all open source parameters.

To modify or add configuration parameters not displayed on the page, use the configuration item customization function. The new parameters will be saved in the component's configuration file and take effect after the component is restarted.

### Prerequisites

- You have understood the meanings of parameters to be added, configuration files that have taken effect, and the impact on components.
- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Impact on the System

- After configuring properties of a service, restart the service if the service status is **Expired**. During the restart, the service cannot be accessed.
- After configuring the HBase, HDFS, Hive, Spark, YARN and MapReduce service properties, download and install the client again or download the configuration file to update the client.

## Adding Custom Parameters on the Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.


**Step 3** On the MRS cluster details page, click **Components**.

**Step 4** Select the target service from the service list.

**Step 5** Click **Service Configuration**.

**Step 6** In the configuration type drop-down box on the right side, switch **Basic** to **All**.





**Figure 6-30** All configurations



**Step 7** In the navigation tree on the left, locate a level-1 node and select **Customization**. The system displays the customized parameters of the current component.

- You can customize parameters for services and roles as required.
- Adding custom parameters for a single role instance is not supported.
- The configuration files that save the newly added customized parameters are displayed in the **Parameter File** column. Each configuration file may support open source parameters with the same name. The effective result is determined by the order in which the component loads the configuration files.

**Step 8** Based on the configuration files and parameter functions, locate the row where a specified parameter resides, enter the parameter name supported by the component in the **Parameter** column and enter the parameter value in the **Value** column.

- You can click  or  to add or delete a custom parameter. You can delete a customized parameter only after you click  for the first time.
- If you want to cancel the modification of a parameter value, click  to restore it.

**Step 9** Click **Save Configuration** and operate as prompted.

----End

## Adding Custom Parameters on Manager

**For MRS 3.x and later versions:**

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Cluster > Services**.

**Step 3** Click the specified service name on the service management page.

**Step 4** Choose **Configurations > All Configurations**.

**Step 5** In the navigation tree on the left, locate a level-1 node and select **Customization**. The system displays the customized parameters of the current component.

- You can customize parameters for services and roles as required.
- Adding custom parameters for a single role instance is not supported.
- The configuration files that save the newly added customized parameters are displayed in the **Parameter File** column. Each configuration file may support open source parameters with the same name. The effective result is determined by the order in which the component loads the configuration files.



**Step 6** Locate the row where a specified parameter resides, enter the parameter name supported by the component in the **Name** column and enter the parameter value in the **Value** column.

You can click + or - to add or delete a customized parameter.

**Step 7** Click **Save**. In the displayed **Save Configuration** dialog box, confirm the modification and click **OK**. After the system displays "Operation succeeded", click **Finish**. The configuration is saved.

Restart the expired service or instance for the configuration to take effect.

 **NOTE**

If  is displayed before a parameter, this parameter takes effect dynamically. After the configuration is saved, the parameter value is automatically updated to the configuration file.  is supported only in MRS 3.2.0 or later.

----End

**For MRS 2.x and earlier:**

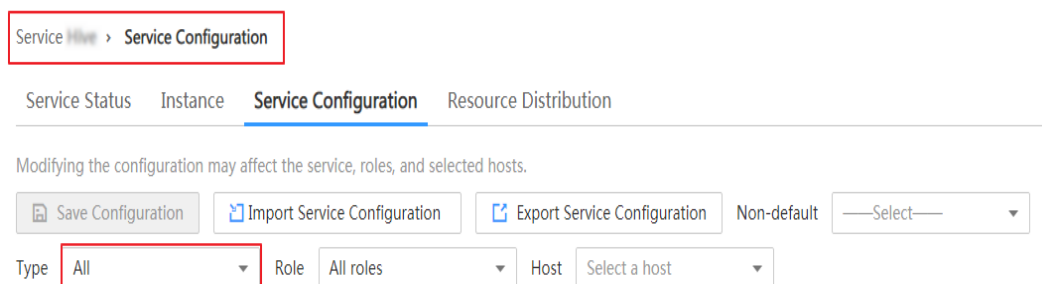
**Step 1** On MRS Manager, click **Services**.

**Step 2** Select the target service from the service list.

**Step 3** Click **Service Configuration**.

**Step 4** Set **Type** to **All**.


**Figure 6-31** Service configurations







**Step 5** In the navigation tree, select **Customization**. The custom parameters of the current component are displayed on Manager.

- You can customize parameters for services and roles as required.
- Adding custom parameters for a single role instance is not supported.
- The configuration files that save the newly added customized parameters are displayed in the **Parameter File** column. Each configuration file may support open source parameters with the same name. The effective result is determined by the order in which the component loads the configuration files.

**Figure 6-32** Custom service configurations



**Step 6** Based on the configuration files and parameter functions, locate the row where a specified parameter resides, enter the parameter name supported by the component in the **Name** column and enter the parameter value in the **Value** column.

- You can click  or  to add or delete a user-defined parameter. You can delete a custom parameter only after you click  for the first time.
- Click  to cancel the change to a parameter value.

**Step 7** Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK** to restart the services.

After the system displays "Operation succeeded", click **Finish**. The service is started.

----End

## Task Example: Adding Custom Hive Parameters

Hive depends on HDFS. By default, Hive accesses the HDFS client. The configuration parameters to take effect are controlled by HDFS in a unified manner.

For example, the HDFS parameter **ipc.client.rpc.timeout** affects the RPC timeout period for all clients to connect to the HDFS server. If you need to modify the timeout period for Hive to connect to HDFS, you can use the configuration customization function. After this parameter is added to the **core-site.xml** file of Hive, this parameter can be identified by the Hive service and its configuration overwrites the parameter configuration in HDFS.

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS cluster details page, click **Components**.


**Step 4** Choose **Hive > Service Configuration**.

**Step 5** In the configuration type drop-down box on the right side, switch **Basic** to **All**.

**Step 6** In the navigation tree on the left, select **Customization** for the Hive service. The system displays the customized service parameters supported by Hive.

**Step 7** In **core-site.xml**, locate the row that contains the **core.site.customized.configs** parameter, enter **ipc.client.rpc.timeout** in the **Parameter** column, and enter a new value in the **Value** column, for example, **150000**. The unit is millisecond.


**Figure 6-33** Configuring custom Hive parameters



**Step 8** Click **Save Configuration** and operate as prompted.



**Figure 6-34** Saving custom configurations



----End

## 6.5.9 Managing MRS Role Instances

You can start a role instance that is in the **Stopped**, **Failed to stop** or **Failed to start** status, stop an unused or abnormal role instance or restart an abnormal role instance to recover its functions.

### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Managing Role Instances on the Console

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** On the MRS cluster details page, click **Components**.

**Figure 6-35** Components page

Dashboard Monitor Nodes **Components** Alarms Files Jobs Tenants

You can manage all services installed in the cluster. [Learn more](#)

| Name      | Version    | Operating Status | Health Status |
|-----------|------------|------------------|---------------|
| ▼ Hadoop  | 3.1.1      |                  |               |
| Spark2x   | 2.4.5      | ✔ Started        | ✘ Faulty      |
| HBase     | 2.2.3      | ✔ Started        | ✔ Good        |
| Hive      | 3.1.0      | ✔ Started        | ✘ Faulty      |
| Hue       | 4.7.0      | ✔ Started        | ✔ Good        |
| Kafka     | 2.11-2.4.0 | ✔ Started        | ✔ Good        |
| Flume     | 1.9.0      | ✔ Started        | ✔ Good        |
| Flink     | 1.12.0     | ✔ Started        | ✔ Good        |
| Oozie     | 5.1.0      | ✔ Started        | ✘ Faulty      |
| ZooKeeper | 3.5.6      | ✔ Started        | ✔ Good        |

- Step 4** Select the target service from the service list.
- Step 5** Click the **Instances** tab.
- Step 6** Select the check box on the left of the target role instance.
- Step 7** Click **More**, select operations such as **Start Instance**, **Stop Instance**, **Restart Instance**, **Rolling-restart Instance**, or **Delete Instance** based on site requirements.

[Component Restart Reference Information](#) outlines which instances support or not support rolling restart.

----End

## Managing Role Instances on Manager

- Step 1** Log in to Manager.
- Step 2** Go to the service instance page.
  - MRS 3.x and later versions: Choose **Cluster** > Services, click the name of the target service, and click **Instances**.
  - MRS 2.x and earlier versions: Click **Services** and click **Instances**.
- Step 3** Select the check box on the left of the target role instance.
- Step 4** Choose **More** > **Start Instance**, **Stop Instance**, or **Restart Instance** as prompted.

[Component Restart Reference Information](#) outlines which instances support or not support rolling restart.

----End

## Exporting and Importing Role Instance Configurations

**For clusters of MRS 3.x and later**

**Step 1** Log in to Manager.

**Step 2** Choose **Cluster > Services** and click the name of the service to be operated in the service view.

**Step 3** Click **Instances** and click the instance to be operated.

**Step 4** Click **Instance Configuration** and click **Export** to export the configuration parameter file to the local PC.

**Step 5** On the **Instance Configurations** page, click **Import**, select the configuration parameter file of the instance, and import the file.

----End

**For MRS clusters of 2.x and earlier:**

**Step 1** Log in to MRS Manager.

**Step 2** Click **Services**.

**Step 3** Select a service.

**Step 4** Select a role instance or click **Instances**.

**Step 5** Select a role instance on a specified host.

**Step 6** Click **Instance Configuration**.

**Step 7** Click **Export Instance Configuration** to export the configuration data of a specified role instance, and choose a path for saving the configuration file.

**Step 8** Click **Import Instance Configuration** to import the configuration data of the specified role instance.

Click **Save Configuration** and select **Restart the role instance**. Click **OK**.

After the system displays "Operation succeeded", click **Finish**. The role instance is started.

----End

### 6.5.10 Managing MRS Role Instance Groups

With MRS, you can easily manage instance groups by grouping multiple instances based on a specified principle, such as nodes with the same hardware configuration.

Any changes made to the configuration parameters of an instance group will apply to all instances within that group.

**NOTE**

This section applies to MRS 3.x or later.

### Prerequisites


You have logged in to MRS Manager. For details, see [Accessing MRS Manager](#).

### Managing MRS Role Instance Groups


**Step 1** On Manager, choose **Cluster > Services**.

**Step 2** Click the specified service name on the service management page.

**Step 3** On the displayed page, click **Instance Groups**.

Click  and configure parameters as prompted.

**Figure 6-36** Creating an instance group



**Table 6-22** Instance group configuration parameters



| Parameter   | Description                                                                                                                                                                                                                                           |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group Name  | Instance group name. The value can contain only letters, digits, underscores (_), hyphens (-), and spaces. It must start with a letter, digit, underscore (_), or hyphen (-) and cannot ends with a space. It can contain a maximum of 99 characters. |
| Role        | Role to which an instance group belongs.                                                                                                                                                                                                              |
| Copy From   | Instance group where the configuration parameter values are copied from. If the value is null, default values are used.                                                                                                                               |
| Description | Instance group description. It can contain only letters, digits, commas (,), periods (.), underscores (_), spaces, and line breaks, and can contain a maximum of 200 characters.                                                                      |

 **NOTE**

Each instance must belong to only one instance group. When an instance is installed for the first time, it belongs to the instance group *Role name-DEFAULT* by default.

**Step 4** Click **OK**.

**Step 5** You can perform the following operations on an instance group as needed.

- **Modifying an instance group:** On the **Instance Group** page, locate the target instance group, click , and modify parameters as prompted. The default instance group cannot be modified.
- **Deleting an instance group:** Click **Instance Groups**. On the **Instance Groups** page, locate the row that contains the target instance group. Click . In the displayed dialog box, click **OK**. The default instance group cannot be deleted. You can delete unnecessary or unused instance groups. Before deleting an instance group, migrate all instances in the group to other instance groups. The default instance group cannot be deleted.
- **Viewing instance group information:** On the **Instance Groups** page, select a role. On the **Basic** page, view all instances of the instance group.

 **NOTE**

To move an instance from an instance group to another, perform the following operations:

1. Select the instance to be moved and click **Move**.
2. In the displayed dialog box, select an instance group to which the instance to be moved.

During the migration, the configuration of the new instance group is automatically inherited. If the instance configuration is modified before the migration, the configuration of the instance will be used.

3. Click **OK**.

Restart the expired service or instance for the configuration to take effect.

- **Configuring instance group parameters:** On the **Instance Groups** page, select a role in the navigation pane, click **Configuration**, adjust the configuration parameters, and click **Save**. The configuration will take effect on all instances in the instance group.

----End

## 6.5.11 Modifying MRS Role Instance Parameters

Configuration parameters of each role instance can be modified. In the scenario where instances are migrated to a new cluster or the service is redeployed, the cluster administrator can import or export all configuration data of a service to quickly copy configuration results.

Making changes to configuration parameters, exporting instance configurations, or importing instance configurations will not impact other instances.

### Impact on the System

- After modifying the configuration of a role instance, you need to restart the instance if the instance status is **Expired**. The role instance is unavailable during restart.

- You need to download and update the client configuration files after configuring HBase, HDFS, Hive, Spark, YARN, and MapReduce service properties.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

## Modifying Instance Parameters on the Console

**Step 1** Log in to the MRS console and click the name of the target cluster.

**Step 2** On the MRS cluster details page, click **Components**.

**Step 3** Select the target service from the service list.


**Step 4** Click the **Instances** tab.

**Step 5** Click the target role instance from the role instance list.

**Step 6** Click the **Instance Configuration** tab.

**Step 7** Switch **Basic** to **All** from the drop-down list on the right of the page. All configuration parameters of the role instance are displayed in the navigation tree.

**Step 8** In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

If you want to cancel the modification of a parameter value, click  to restore it.

**Step 9** Click **Save Configuration** and operate as prompted.

----End

## Modifying Instance Parameters on Manager

**Step 1** Log in to Manager and choose **Cluster > Services**.

For MRS 2.x and earlier versions, click **Services**.

**Step 2** Click the specified service name and click **Instances**.

**Step 3** Click the specified instance and select **Instance Configuration**.

By default, **Basic Configuration** is displayed. To modify more parameters, click **All Configurations**. All parameter categories supported by the instance are displayed on the **All Configurations** page.

**Step 4** In the navigation tree, select the specified parameter category and change the parameter values on the right.



Enter the name of a parameter in the upper right corner to search for it. Search results are displayed on Manager in real-time.

**Step 5** Click **Save**. In the confirmation dialog box, click **OK**.

For MRS 2.x and earlier versions, click **Save Configuration**, select **Restart the role instance**, and click **OK** to restart the role instance.

Wait until the message "Operation succeeded" is displayed. Click **Finish**. The configuration is modified.

 **NOTE**

- After the configuration parameters of a role instance are modified, you need to restart the instance if the instance status is **Expired**. You can select the expired instance on the **Instances** page and choose **More > Restart Instance**.
- If  is displayed before a parameter, this parameter takes effect dynamically. After the configuration is saved, the parameter value is automatically updated to the configuration file.  is supported only in MRS 3.2.0 or later.

----End

## 6.5.12 Perform an Active/Standby Switchover for MRS Role Instances

Some service roles are deployed in active/standby mode. If the active instance needs to be maintained and cannot provide services, or other maintenance is required, you can manually trigger an active/standby switchover.

### Prerequisites

You have logged in to MRS Manager. For details, see [Accessing MRS Manager](#).

### Procedure

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Cluster > Services**.

For MRS 2.x and earlier versions, click **Services**. The parameters vary depending on the version.

**Step 3** Click the specified service name (such as **HDFS**).

**Step 4** On the service details page, expand the **More** drop-down list and select **Perform Role Instance Switchover**.

For example, click **Perform NameNode Switchover**.

**Step 5** In the displayed dialog box, enter the password of the current login user and click **OK**.

**Step 6** In the dialog box that is displayed, confirm the operation impact and click **OK** to perform an active/standby switchover for role instances.

 NOTE

- The following components support active/standby switchover of role instances: DBService, HDFS, YARN, Storm, HBase, MapReduce, and Loader. This function cannot be used for other role instances.
- When an active/standby switchover is performed for a NameNode on HDFS, a NameService must be set.

----End

## 6.5.13 Decommissioning and Recommissioning an MRS Role Instance

Data in an MRS cluster can be spread out over multiple Core nodes. If one of these nodes fails, it can destabilize the entire cluster. To prevent this, you can decommission a specific role instance on MRS, which will halt its service provision.

After fault rectification, you can recommission the role instance.

The following role instances can be decommissioned or recommissioned:

- DataNode role instance on HDFS
- NodeManager role instance on YARN
- RegionServer role instance on HBase
- ClickHouseServer role instance of ClickHouse (supported in MRS 3.1.2 or later)
- IoTDBServer role instance of IoTDB
- Broker role instance on Kafka

### Constraints

- If a role instance is out of service, you must recommission the instance to start it before using it again.
- For details about ClickHouseServer instance decommissioning restrictions, see [Constraints on ClickHouseServer Scale-in](#).
- By default, if the number of the DataNodes is less than or equal to that of HDFS replicas, decommissioning cannot be performed. If there are only three HDFS replicas and less than four DataNodes in the system, decommissioning cannot be carried out. In such a scenario, an error will be reported, and Manager will exit the decommissioning process 30 minutes after attempting to perform it.
- During MapReduce task execution, files with 10 replicas are generated. Therefore, if the number of DataNode instances is less than 10, decommissioning cannot be performed.
- If the number of DataNode racks (the number of racks is determined by the number of racks configured for each DataNode) is greater than 1 before the decommissioning, and after some DataNodes are decommissioned, that of the remaining DataNodes changes to 1, the decommissioning will fail. Therefore, before decommissioning DataNode instances, you need to evaluate the impact of decommissioning on the number of racks to adjust the DataNodes to be decommissioned.



- If multiple DataNodes are decommissioned at the same time, and each of them stores a large volume of data, the DataNodes may fail to be decommissioned due to timeout. To avoid this problem, it is recommended that one DataNode be decommissioned each time and multiple decommissioning operations be performed.
- If the number of IoTDBServers is less than or equal to the number of region copies configured for the cluster (3 by default), decommissioning cannot be performed.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).
- If the DataNode is to be decommissioned, perform the following steps to perform a health check before decommissioning:
  - a. Log in to the client installation node as a client user and switch to the client installation directory.
  - b. For a security cluster, use user **hdfs** for permission authentication.

```
source bigdata_env #Configure client environment variables.
kinit hdfs #Configure kinit authentication.
Password for hdfs@HADOOP.COM: #Enter the login password of user hdfs.
```
  - c. Run the **hdfs fsck / -list-corruptfileblocks** command, and check the returned result.
    - If "has 0 CORRUPT files" is displayed, the health check is successful.
    - If the command output does not contain "has 0 CORRUPT files" and the name of the damaged file is returned, perform the following operations to delete the damaged file:

**hdfs dfs -rm** *Name of the damaged file*


### NOTE

Deleting a file is a high-risk operation. Ensure that the files are no longer needed before performing this operation.

## Decommissioning or Recommissioning an Instance on the Console

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** On the MRS cluster details page, click **Components**.
- Step 4** Click a service in the service list.
- Step 5** Click the **Instances** tab.
- Step 6** Select an instance.
- Step 7** Choose **More > Decommission** or **Recommission** to perform the corresponding operation.

Figure 6-37 Decommissioning an instance



| Host Name        | OM IP Address | Business IP Address | Rack              | Operating Status | Health Status | Configuration Status |
|------------------|---------------|---------------------|-------------------|------------------|---------------|----------------------|
| node-master1m0FW | 192.168.0.187 | 192.168.0.187       | /default/rack1485 | Started          | Good          | Synchronized         |
| node-master2m0QR | 192.168.0.82  | 192.168.0.82        | /default/rack1485 | Started          | Good          | Synchronized         |

**NOTE**

During the instance decommissioning, if the service corresponding to the instance is restarted in the cluster using another browser, MRS displays a message indicating that the instance decommissioning is stopped, but the **Operating Status** of the instance is displayed as **Started**. In this case, the instance has been decommissioned on the background. You need to decommission the instance again to synchronize the operating status.

----End

## Decommissioning or Recommissioning Instances on Manager

**Step 1** Log in to Manager and go to the page of the component instance to be operated.

- MRS 3.x and later versions: Choose **Cluster** > **Services**, click the specified service name on the **Services** page, and click **Instances**.
- MRS 2.x and earlier versions: Click **Services**, click the specified service name in the service list, and click **Instances**.

**Step 2** Select the specified role instance to be decommissioned.

**Step 3** Select **Decommission** or **Recommission** from the **More** drop-down list.

In the displayed dialog box, enter the password of the current login user and click **OK**.

Select the operation impact and click **OK** to perform the corresponding operation.

**NOTE**

During the instance decommissioning, if the service corresponding to the instance is restarted in the cluster using another browser, MRS displays a message indicating that the instance decommissioning is stopped, but the **Operating Status** of the instance is displayed as **Started**. In this case, the instance has been decommissioned on the background. You need to decommission the instance again to synchronize the operating status.

----End

### 6.5.14 Enabling and Disabling Ranger Authentication for an MRS Component

By default, the Ranger service is installed and Ranger authentication is enabled for a newly installed cluster with Kerberos authentication enabled. You can use the permission plug-in of the component to establish fine-grained security access policies for accessing component resources. If Ranger authentication is not required, the cluster administrator can manually disable Ranger authentication on the service page. After Ranger authentication is disabled, the system continues to perform permission control based on the role model of Manager when accessing component resources.

In a cluster upgraded from an earlier version, Ranger authentication is not used by default when users access component resources. The cluster administrator can manually enable Ranger authentication after installing the Ranger service.

#### NOTE

- This section applies only to MRS 3.x or later.
- In a cluster in security mode, the following components support Ranger authentication: HDFS, YARN, Kafka, Hive, HBase, Storm, Impala, HetuEngine, CDL, and Spark/Spark2x.
- In a cluster in non-security mode, Ranger supports permission control on component resources based on OS users. The following components support Ranger authentication: HBase, HDFS, Hive, Spark/Spark2x, and YARN.
- After Ranger authentication is enabled, all authentication of the component will be managed by Ranger. The permissions set by the original authentication plug-in will become invalid (The ACL rules of HDFS and YARN components still take effect). Exercise caution when performing this operation. You are advised to deploy permissions on Ranger in advance.
- After Ranger authentication is disabled, all authentication of the component will be managed by the permission plug-in of the component. The permission set on Ranger will become invalid. Exercise caution when performing this operation. You are advised to deploy permissions on Manager in advance.

## Enabling Ranger Authentication

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Cluster > Services**.

**Step 3** Click the specified service name on the service management page.

**Step 4** On the service details page, expand the **More** drop-down list and select **Enable Ranger**.

**Step 5** In the displayed dialog box, enter the password of the current login user and click **OK**.

**Step 6** In the service list, restart the service whose configuration has expired.

----End

## Disabling Ranger Authentication

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Cluster > Services**.

**Step 3** Click the specified service name on the service management page.

**Step 4** On the service details page, expand the **More** drop-down list and select **Disable Ranger**.

**Step 5** Enter the password of the current login user and click **OK**. In the displayed dialog box, click **OK**.

**Step 6** In the service list, restart the service whose configuration has expired.

----End

## 6.5.15 Accessing Web Pages of Open Source Components Managed in MRS Clusters

If the component has an open source web UI, you can access it by clicking the web UI link in the component's basic information area.

### NOTE

For clusters with Kerberos authentication enabled, user **admin** does not have the management permission on each component. To access the web UI of each component, create a user who has the management permission on the corresponding component by referring to [Creating an MRS Cluster User](#).

### Accessing the Web UI of an Open Source Component

**Step 1** Log in to MRS Manager.

For details, see [Accessing MRS Manager](#).

**Step 2** [Table 6-23](#) displays the steps to locate the entry for each component.

**Table 6-23** Web UI addresses of open-source components

| Web UI Type                | Address (Earlier than MRS 3.x)                                                                                                                    | Address (MRS 3.x and Later)                                                                                                                                            |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HDFS Name Node             | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; HDFS &gt; NameNode Web UI &gt; NameNode (Active)</b> .               | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; HDFS &gt; NameNode Web UI &gt; NameNode (Host name, Active)</b> .                      |
| HBase HMaster              | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; HBase &gt; HMaster Web UI &gt; HMaster (Active)</b> .                | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; HBase &gt; HMaster Web UI &gt; HMaster (Host name, Active)</b> .                       |
| MapReduce JobHistoryServer | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; MapReduce &gt; JobHistoryServer Web UI &gt; JobHistoryServer</b> .   | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; MapReduce &gt; JobHistoryServer Web UI &gt; JobHistoryServer (Host name, Active)</b> . |
| YARN ResourceManager       | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; Yarn &gt; ResourceManager Web UI &gt; ResourceManager (Active)</b> . | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; Yarn &gt; ResourceManager Web UI &gt; ResourceManager (Host name, Active)</b> .        |

| Web UI Type      | Address (Earlier than MRS 3.x)                                                                                                                                                                                                                                         | Address (MRS 3.x and Later)                                                                                                                                                                                                                                                          |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Spark JobHistory | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; Spark &gt; Spark Web UI &gt; JobHistory</b> .                                                                                                                                             | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; Spark2x &gt; Spark2x Web UI &gt; JobHistory2x (Host name, Active)</b> .                                                                                                                              |
| Hue              | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; Hue &gt; Hue Web UI &gt; Hue (Active)</b> .<br>Loader is a graphical data migration management tool based on the open-source Sqoop web UI, and its interface is hosted on the Hue web UI. | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; Hue &gt; Hue Web UI &gt; Hue (Host name, Active)</b> .<br>Loader is a graphical data migration management tool based on the open-source Sqoop web UI, and its interface is hosted on the Hue web UI. |
| Tez              | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; Tez &gt; Tez Web UI &gt; TezUI</b> .                                                                                                                                                      | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; Tez &gt; Tez Web UI &gt; TezUI (Host name)</b> .                                                                                                                                                     |
| Presto           | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; Presto &gt; Presto Web UI &gt; Coordinator (Active)</b> .                                                                                                                                 | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; Presto &gt; Coordinator Web UI &gt; Coordinator(Coordinator)</b> .                                                                                                                                   |
| Ranger           | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; Ranger &gt; Ranger Web UI &gt; RangerAdmin (Active)</b> .                                                                                                                                 | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; Ranger &gt; Ranger Web UI &gt; RangerAdmin</b> .                                                                                                                                                     |
| Storm            | Earlier than MRS 3.x: On the cluster details page, choose <b>Components &gt; Storm &gt; Storm Web UI &gt; UI</b> .                                                                                                                                                     | MRS 3.x or later: On the Manager homepage, choose <b>Cluster &gt; Services &gt; Storm &gt; Storm Web UI &gt; UI (Host name)</b> .                                                                                                                                                    |

----End

## 6.6 Managing MRS Cluster Nodes

### 6.6.1 Checking the Running Status of an MRS Cluster Node

You can keep track of the status of each node in a running MRS cluster in real-time on the console or Manager. This helps you promptly detect any potential resource problems.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

## Checking the Node Status on the Console

**Step 1** Log in to the MRS console.


**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the cluster details page, click **Nodes**.

**Step 4** Expand the node group name to view basic information about the nodes in a node group, such as their status, CPU usage, memory usage, and disk usage.

In the host list, click a host name to view detailed node information.

**Table 6-24** MRS cluster node information

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Group   | Node group name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Node Type    | Node type corresponding to the node group. For details, see <a href="#">MRS Cluster Node Types</a> .<br><br>On the <b>Nodes</b> page, click  next to a node group name to unfold the nodes contained in the node group. Click a node name to remotely log in to the ECS using the password or key pair configured during cluster creation. For details about the parameters, see <a href="#">Viewing MRS Cluster Component Monitoring Metrics</a> . |
| Node Count   | Number of nodes in a node group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Billing Mode | Billing mode of the cluster you purchased, including <b>Pay-Per-Use</b> and <b>Yearly/Monthly</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                    |

----End

## Checking the Node Status on Manager (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Click **Hosts** to view the host list.

The host list shows all the hosts in the cluster along with their basic information. You can use this list to perform O&M operations on hosts, apply filters, and search for specific hosts. Additionally, you can click **Export All**, set **Save As** to **TXT** or **CSV**, and click **OK** to export information about all hosts. **Host View** is displayed by

default. You can click **Role View** to switch the view type. You can also click the edit button to customize the content displayed in each view.

**Table 6-25** View type description

| View Type | Description                                                                                                                                |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------|
| Host View | Displays the IP address, rack planning, running status, and hardware resource usage of each host.                                          |
| Role view | Displays the roles that have been deployed on each host. If the role supports the active/standby mode, the role name is displayed in bold. |

**Table 6-26** Host running status description

| Status    | Description                                                       |
|-----------|-------------------------------------------------------------------|
| Normal    | Indicates that the host is in the normal state.                   |
| Faulty    | Indicates that the host is abnormal.                              |
| Unknown   | Indicates that the initial status of the host cannot be detected. |
| Isolated  | Indicates that the host is isolated.                              |
| Suspended | Indicates that the host is stopped.                               |

**Step 3** In the host list, click the specified host name to view the host overview information.

The host details page contains the basic information area, disk status area, instance list area, and monitoring graphs.

**Table 6-27** Host details

| Item                   | Description                                                                                                                                                                                                            |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Basic Information      | The basic information area contains the key information about the host, such as the management IP address, service IP address, host type, rack, firewall, number of CPU cores, and OS.                                 |
| Disk                   | The disk status area contains all disk partitions configured for the cluster on the host and the usage of each disk partition.                                                                                         |
| Instances              | The instance list displays all role instances installed on the host and the status of each role instance. You can click the log file next to a role instance name to view the log file content of the instance online. |
| Alarm and Event Record | The alarm and event history area displays the key alarms and events reported by the current host. The system can display a maximum of 20 historical records.                                                           |

| Item  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Chart | <p>The monitoring chart area is displayed on the right of the host details page, and contains the key monitoring metrics of the host.</p> <p>You can choose ▾ &gt; <b>Customize</b> in the upper right corner to customize the monitoring reports to be displayed in the chart area. Select a time range and choose ▾ &gt; <b>Export</b> to export detailed monitoring metric data within the specified time range.</p> <p>You can click ⓘ next to the title of a monitoring indicator to open the description of the monitoring metric.</p> |

**Step 4** View the host chart, process, and resource information.

Click the **Chart**, **Process**, and **Resource** to view the full monitoring chart information about the host.

- The **Chart** page displays all monitoring graphs of the host.
- The **Process** page shows information about the role processes of the deployed service instances on the current host, including their status, PID, and running time. You can easily access the log files of each process online.
- The **Resource** page displays the resource usage details of the deployed service instances on the current host, such as CPU, memory, disk, and port usage.

----End

## Checking the Node Status on Manager (MRS 2.x or Earlier)

**Step 1** Log in to MRS Manager.

**Step 2** Click **Hosts** to view the status of all hosts.

The following table lists the operating status and health status of the host.

**Table 6-28** Host operating status

| Status   | Description                                                           |
|----------|-----------------------------------------------------------------------|
| Normal   | The host and service roles on the host are running properly.          |
| Isolated | The host is isolated, and the service roles on the host stop running. |

**Table 6-29** Host health status

| Status | Description                                       |
|--------|---------------------------------------------------|
| Good   | The host can properly send heartbeats.            |
| Bad    | The host fails to send heartbeats due to timeout. |



| Status  | Description                                        |
|---------|----------------------------------------------------|
| Unknown | Initial status of the host when the host is added. |

**Step 3** Click the target host in the host list to view its status and metric information.

**Step 4** Customize and export monitoring charts.

1. In the **Charts** area, click **Customize** to customize service monitoring metrics.
2. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.
3. Click **Export** to export the displayed metrics.

----End

## 6.6.2 Starting and Stopping All Roles on an MRS Cluster Node

When a host or node is faulty, you may need to stop all roles on the host on MRS to check the host. After the fault is rectified, start all roles on the host to recover host services.

### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Starting or Stopping a Node Role on the Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS details page, click **Nodes**.

**Step 4** Unfold the node group information and select the check box of the target node.

**Step 5** Choose **Node Operation** > **Start All Roles** or **Stop All Roles** to perform the required operation.

----End

### Starting or Stopping Node Roles on Manager

**Step 1** Log in to the cluster Manager.

**Step 2** Click **Hosts** to go to the host list page.

For clusters of MRS 2.x or earlier, click **Hosts**.

**Step 3** Select the check box of the target hosts.

**Step 4** Select **Start All Instances** or **Stop All Instances** from the **More** drop-down list to start or stop all role instances.

----End

### 6.6.3 Isolating an MRS Cluster Node

If a host is found to be abnormal or faulty, affecting cluster performance or preventing services from being provided, you can temporarily exclude that host from the available nodes in the cluster. In this way, the client can access other available nodes. In scenarios where patches are to be installed in a cluster, you can also exclude a specified node from patch installation.

Only non-management nodes can be isolated.

You can isolate a host manually on MRS based on the actual service requirements or O&M plan.

#### Impact on the System

- After a host is isolated, all role instances on the host will be stopped. You cannot start, stop, or configure the host and any instances on the host.
- After a host is isolated, statistics of the monitoring status and indicator data of the host hardware and instances cannot be collected or displayed.
- For some services, after a host is isolated, some instances on other nodes do not work, and the service configuration status may expire.
- Retain the default SSH port (22) of the target node. Otherwise, the task described in this section will fail.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

#### Isolating a Host on the Console

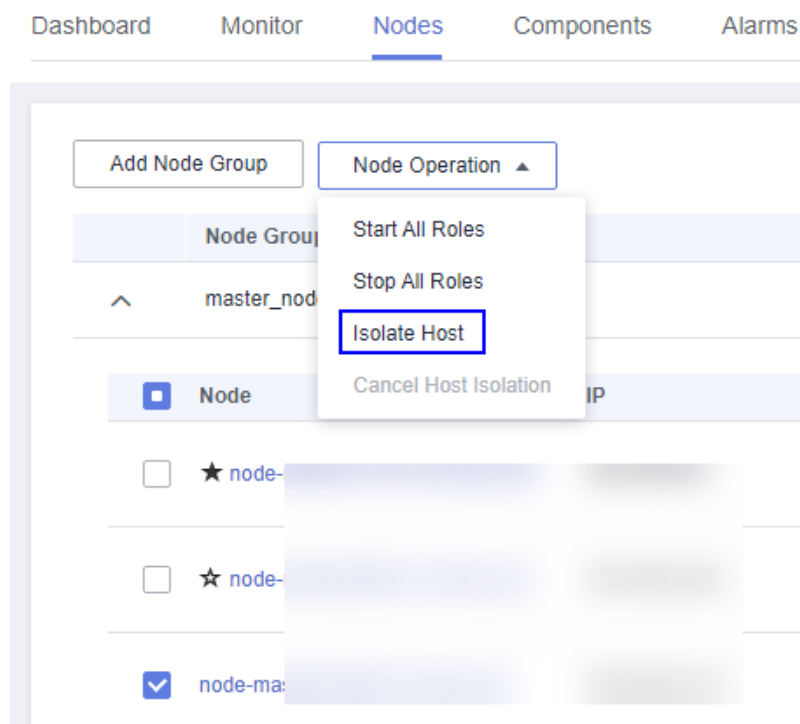
**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS details page, click **Nodes**.

**Step 4** Unfold the node group information and select the check box of the target host.

**Step 5** Choose **Node Operation** > **Isolate Host**.

**Figure 6-38** Isolating a host

**Step 6** Confirm the information about the host to be isolated and click **OK**.

When **Operation successful** is displayed, click **Finish**. The host is isolated successfully, and the value of **Operating Status** becomes **Isolated**.

**NOTE**

The isolation of a host can be canceled and the host can be added to the cluster again. After the exception or fault of a host is handled, you must cancel the isolation of the host for proper usage.

**Step 7** Cancel the isolation status of the host before using the host if you have rectified the host exception or fault.

1. On the **Nodes** page, expand the node group information.
2. Select the check box of the host to be de-isolated and choose **Node Operation > Cancel Host Isolation**.
3. Confirm the information about the host for which the isolation is to be cancelled and click **OK**.

When **Operation successful** is displayed, click **Finish**. The host is de-isolated, and the value of **Operating Status** becomes **Normal**.

----End

## Isolating a Host on Manager

For MRS 3.x or later:

**Step 1** Log in to FusionInsight Manager and click **Hosts**.

**Step 2** Select the check box of the host to be isolated, choose **More > Isolate Host**, enter the password for identity authentication, and click **OK**.

**Step 3** Confirm the information about the host to be isolated and click **OK**.

After the system displays a message indicating that the operation is successful, click **Finish**. If the host is isolated, its **Running Status** is displayed as **Isolated**.

**Step 4** Log in to the isolated host as user **root** and run the **ps -ef | grep 'container' | grep '\${BIGDATA\_HOME}' | awk '{print \$2}' | xargs -l '{}' kill -9 '{}'** command to find and stop the container process.

**Step 5** Cancel the isolation status of the host before using the host if you have rectified the host exception or fault.

On the **Hosts** page, select the isolated host and choose **More > Cancel Isolation**.

 **NOTE**

After the isolation is canceled, all role instances on the host are not started by default. To start role instances on the host, select the target host on the **Hosts** page and choose **More > Start All Instances**.

----End

For MRS 2.x or earlier:

**Step 1** Log in to MRS Manager, and choose **Hosts**.

**Step 2** Select the check box of the host to be isolated, choose **More > Isolate Host**, enter the password for identity authentication, and click **OK**.

**Step 3** Confirm the information about the host to be isolated and click **OK**.

After the system displays a message indicating that the operation is successful, click **Finish**. If the host is isolated, its **Operating Status** is displayed as **Isolated**.

**Step 4** Cancel the isolation status of the host before using the host if you have rectified the host exception or fault.

1. On MRS Manager, click **Hosts**.
2. Select the check box of the host to be de-isolated and choose **More > Cancel Host Isolation**.
3. Click **OK** in the displayed dialog box.  
When **Operation successful** is displayed, click **Finish**. The host is de-isolated successfully, and the value of **Operating Status** becomes **Normal**.
4. Click the name of the de-isolated host to show its status, and click **Start All Roles**.

----End

## 6.6.4 Modifying the Rack Information of an MRS Cluster Node

All hosts in a large cluster are usually deployed on multiple racks. Hosts on different racks communicate with each other through switches. The network bandwidth between different hosts on the same rack is much greater than that on

different racks. In this case, plan the network topology based on the following requirements:

- To improve the communication speed, it is recommended that data be exchanged between hosts on the same rack.
- To improve the fault tolerance capability, distribute processes or data of distributed services on different hosts of multiple racks as dispersedly as possible.

Hadoop uses a file directory structure to represent hosts.

The HDFS cannot automatically determine the network topology of each DataNode in the cluster. You need to set the rack name to identify the rack where the host is located so that the NameNode can draw the network topology of the required DataNodes and back up data of the DataNodes to different racks. Similarly, YARN needs to obtain rack information and allocate tasks to different NodeManagers as required.

If the cluster network topology changes, you need to reallocate racks for hosts on FusionInsight Manager so that related services can be automatically adjusted.

#### NOTE

This section applies only to MRS 3.x or later.

## Impact on the System

- If the name of the host rack is changed, storage policy for HDFS replicas, YARN task assignment, and storage location of Kafka partitions will be affected. After the modification, you need to restart the HDFS, YARN, and Kafka for the configuration to take effect.
- Improper rack configuration will unbalance loads (including CPU, memory, disk, and network) among nodes in the cluster, which decreases the cluster reliability and stability. Therefore, before allocating racks, take all aspects into consideration and properly set racks.

## Rack Allocation Policies

#### NOTE

Physical rack: indicates the real rack where the host resides.

Logical rack: indicates the rack name of the host on FusionInsight Manager.

Policy 1: Each logical rack has nearly the same number of hosts.

Policy 2: The name of the logical rack of the host must comply with that of the physical rack to which the host belongs.

Policy 3: If there are only few hosts on a physical rack, combine this physical rack and other physical racks with few hosts into a logical rack, which complies with policy 1. Hosts in two equipment rooms cannot be placed in one logical rack. Otherwise, performance problems may be caused.

Policy 4: If there are lots of hosts on a physical rack, divide these hosts into multiple logical racks, which complies with policy 1. Hosts with great differences should not be placed in the same logical rack. Otherwise, the cluster reliability will be decreased.

Policy 5: You are advised to set **default** or other values for logical racks on the first layer, and the values in the same cluster must be consistent.

Policy 6: The number of hosts in each rack cannot be less than 3.

Policy 7: A cluster can contain at most 50 logical racks. If there are too many logical racks in a cluster, the maintenance is difficult.

## Best Practices

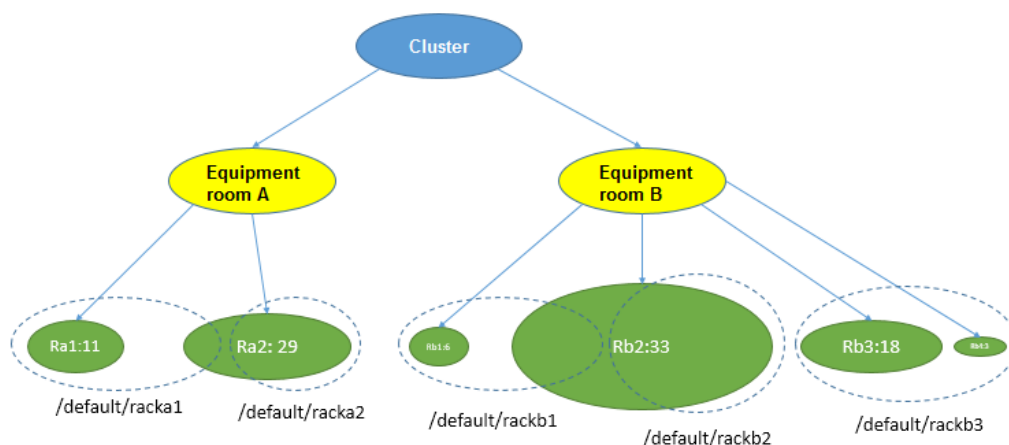
For example, in a cluster, 100 hosts are located in two equipment rooms A and B. A has 40 hosts and B has 60 hosts. In room A, there are 11 hosts on physical rack Ra1 and 29 hosts on physical rack Ra2. In room B, there are six hosts on physical rack Rb1, 33 hosts on physical rack Rb2, 18 hosts on physical rack Rb3, and three hosts on physical rack Rb4.

According to the rack allocation policy, each logical rack contains nearly the same number (for example, 20) of hosts. The allocation details are as follows:

- Logical rack /default/racka1: 11 hosts on physical rack Ra1 and nine hosts on physical rack Ra2
- Logical rack /default/racka2: the remaining 20 hosts (except the nine hosts of logical rack /default/racka1) on physical rack Ra2
- Logical rack /default/rackb1: six hosts on physical rack Rb1 and 13 hosts on physical rack Rb2
- Logical rack /default/rackb2: the remaining 20 hosts on physical rack Rb2
- Logical rack /default/rackb3: 18 hosts on physical rack Rb3 and three hosts on physical rack Rb4

Rack allocation example:

**Figure 6-39** Rack division



## Configuring the Cluster Node Rack

**Step 1** Log in to FusionInsight Manager.

**Step 2** Click **Hosts**.

**Step 3** Select the check box of the target host.

**Step 4** Select **Set Rack** from the **More** drop-down list.

- Set rack names in hierarchy based on the actual network topology. Separate racks from different layers using slashes (/).
- Rack naming rules are as follows: */level1/level2/...*. The number of levels must be at least 1, and the name cannot be empty. A rack can contain letters, digits, and underscores (\_) and cannot exceed 200 characters.  
For example, */default/rack0*.
- If the hosts in the rack to be modified contain DataNode instances, ensure that the rack name levels of the hosts where all DataNode instances reside are the same. Otherwise, the configuration fails to be delivered.

**Step 5** Click **OK**.

----End

## 6.6.5 Scaling Up Master Node Specifications in an MRS Cluster

As users' increasing services lead to Core node scale-out and high CPU usage, Master node specifications cannot meet user requirements and need to be scaled up. This section describes how to scale up Master node specifications.

### Prerequisites

- You have checked whether the Host Security Service (HSS) is enabled. If HSS is enabled, disable the HSS monitoring on the MRS cluster before you scale up master node specifications.
- Ensure that sufficient specification resources are available throughout the steps in [Scaling Up Master Node Specifications \(Step-by-Step Upgrade\)](#).

### Use Restrictions

- Master nodes can be scaled up for clusters with two or more master nodes.
- The specifications of the Master node in a BMS cluster cannot be upgraded.
- For MRS 1.8.2 or later to a version earlier than MRS 3.x or MRS 3.1.0 or later, see [Scaling Up Master Node Specifications \(One-Click Upgrade\)](#).
- For MRS 3.0.5 and a version earlier than MRS 1.8.2, see [Scaling Up Master Node Specifications \(Step-by-Step Upgrade\)](#).
- Do not perform other operations on the cluster during the scale-up.
- It is best to scale up the Master node specifications during off-peak hours to prevent any service interruptions.

## Scaling Up Master Node Specifications (One-Click Upgrade)

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the **Nodes** tab page, select **Scale Up Specifications** in the **Operation** column of the Master node group. The **Scale Up Master Node Specifications** page is displayed.

**Step 4** Select the target specifications and click **Submit Order**. The order has been submitted successfully.

The node specification scale-up takes some time. After the scale-up is successful, the cluster status changes to **Running**.

 **NOTE**

- The VM to be scaled up is automatically stopped during the scale-up and started after the scale-up is complete.
- The scale-up does not automatically upgrade the memory of components due to different component usage requirements. You can adjust the memory of components as needed.

----End

## Scaling Up Master Node Specifications (Step-by-Step Upgrade)

### Preparing for Scaling Up Master Node Specifications

**Step 1** Log in to the MRS console.

**Step 2** In the navigation pane on the left, choose **Active Clusters**, select the cluster for which you want to scale up Master node specifications, and click its name to switch to the cluster details page.

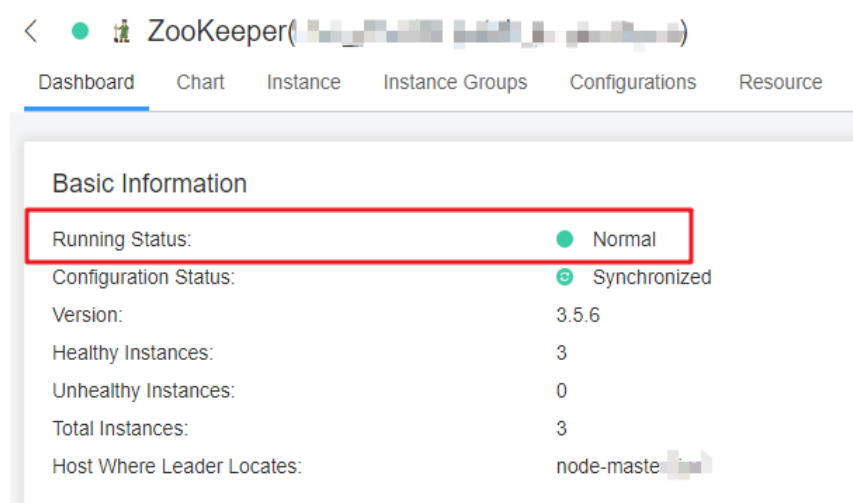
**Step 3** Ensure that the cluster status is **Running**.

**Step 4** On the **Nodes** tab page, ensure that all nodes in the cluster are in the **Running** state.

**Step 5** Log in to Manager and go to the cluster management page. For details, see [Accessing MRS Manager](#).

**Step 6** Choose **Cluster > Services > ZooKeeper > Dashboard** and ensure that **Running Status** of the ZooKeeper service is **Normal**.

**Figure 6-40** ZooKeeper service status



**Step 7** Update service parameter settings as required. For details, see [Modifying the Configuration Parameters of an MRS Cluster Component](#).




 **NOTE**

You need to perform this step only once before scaling up the standby Master node.

**Step 8** Choose **Cluster > Services > HDFS > Instance**.

**Step 9** Record the business IP address of **NameNode (Standby)**. When upgrading the specifications of the active Master node, record the business IP address of **NameNode (Active)**. **Figure 6-41** shows the location of the business IP address.


**Figure 6-41** Business IP address of the NameNode



| Role                        | Running Status | Configuration Status | Host Name       | Management IP Address | Service IP Address | Rack              | Instance Group   |
|-----------------------------|----------------|----------------------|-----------------|-----------------------|--------------------|-------------------|------------------|
| Zkfc                        | Normal         | Synchronized         | node-master2e0f | 192.168.1.100         | 192.168.1.100      | /default rack5501 | Zkfc-DEFAULT     |
| Zkfc                        | Normal         | Synchronized         | node-master0y4  | 192.168.1.101         | 192.168.1.101      | /default rack5501 | Zkfc-DEFAULT     |
| NameNode (hacuster Standby) | Normal         | Synchronized         | node-master2e0f | 192.168.1.100         | 192.168.1.100      | /default rack5501 | NameNode-DEFAULT |
| NameNode (hacuster Active)  | Normal         | Synchronized         | node-master0y4  | 192.168.1.101         | 192.168.1.101      | /default rack5501 | NameNode-DEFAULT |

 **NOTE**

Only when the cluster is an analysis cluster, you can perform **Step 8** to **Step 9** to record the IP addresses of the active and standby nodes.

**Step 10** On the upper right of the Manager page, check the number next to the  icon. If the number is 0, there is no running tasks in the cluster.

**Step 11** Click **Hosts**. If the cluster is an analysis cluster, select the checkbox of the host corresponding to the business IP address of the **NameNode** recorded in **Step 9**. If the cluster is a streaming cluster, you do not need to distinguish the active and standby nodes. You only need to choose hosts for the scale-up.

**Step 12** Choose **More > Stop All Instances** and wait until all instances are stopped.

 NOTE

- When the node where Manager resides is scaled up, Manager may not be accessed due to an active/standby switchover. It is a normal phenomenon. Try to log in to Manager later. If the login fails for a long time, contact O&M personnel.
- After all roles are stopped, the following alarms may be generated. After the scale-up of Master node specifications is complete and all roles are started, the alarms are automatically cleared.
  - [ALM-12006 Node Fault](#)
  - [ALM-12010 Manager Heartbeat Interruption Between the Active and Standby Nodes](#)
  - [ALM-12039 Active/Standby OMS Databases Not Synchronized](#)
  - [ALM-14000 HDFS Service Unavailable](#)
  - [ALM-14010 NameService Service Is Abnormal](#)
  - [ALM-14012 JournalNode Is Out of Synchronization](#)
  - [ALM-16004 Hive Service Unavailable](#)
  - [ALM-18000 Yarn Service Unavailable](#)
  - [ALM-19000 HBase Service Unavailable](#)
  - [ALM-20002 Hue Service Unavailable](#)
  - [ALM-27001 DBService Service Unavailable](#)
  - [ALM-27003 DBService Heartbeat Interruption Between the Active and Standby Nodes](#)
  - [ALM-27004 Data Inconsistency Between Active and Standby DBServices](#)
  - [ALM-43001 Spark2x Service Unavailable](#)

----End

### Scaling Up Master Node Specifications

**Step 1** Log in to the MRS console.

**Step 2** In the navigation pane on the left, choose **Active Clusters**, select the cluster for which you want to scale up Master node specifications, and click its name to switch to the cluster details page.

**Step 3** On the **Nodes** tab page, select **Scale Up Specifications** in the **Operation** column of the Master node group.

**Step 4** Select the target specifications and click **Next**.

 NOTE


Ensure that target specification resources are sufficient. Otherwise, the active node cannot be scaled up.

**Step 5** On the **Confirm** page that is displayed, confirm the target node specifications and fees and click **OK**.

**Step 6** Ensure that all services on the standby Master node have been stopped. For operation details, refer to [Step 1](#) to [Step 12](#) in the **Preparing for Scaling Up Master Node Specifications** part. On the **Scale Up Master Node Specifications** page, select **Are you sure you have stopped all services on the standby Master node?** and **If not all services are stopped before the scale-up, data saving failure or data damage may occur.** and click **Submit Order**.



Figure 6-43 NameNode



The screenshot shows the Hadoop Manager interface with a green navigation bar containing 'Hadoop', 'Overview', 'Datanodes', 'Datanode Volume Failures', 'Snapshot', 'Startup Progress', and 'Utilities'. The main content area is titled 'Overview' and shows details for a NameNode in a 'standby' state. A table lists the following information:

|                |                                                        |
|----------------|--------------------------------------------------------|
| Namespace:     | hacluster                                              |
| Namenode ID:   | 30                                                     |
| Started:       | 2024-09-23 10:00:00 +0800 (UTC+8)                      |
| Version:       | 3.6.7-cdh360-UB6 (rhel7)                               |
| Compiled:      | Thu Oct 28 11:00:00 +0800 2021 by jenkins from/centos7 |
| Cluster ID:    | hacluster                                              |
| Block Pool ID: | BP-12345678-123456789-123456789                        |

- Step 10** Stop all services on the active Master node by referring to [Step 1](#) to [Step 12](#) in the **Preparing for Scaling Up Master Node Specifications** part.
- Step 11** On the **Scale Up Master Node Specifications** page, select **I confirm that all services on the standby Master node have been started.** and **I confirm that all services on the active Master node have been stopped,** and click **Submit**.
- Step 12** On the **Confirm** page that is displayed, confirm again that all services on the active Master node are stopped and click **OK** to start scaling up the specifications of the active Master node.
- The node specification scale-up takes some time. Please wait. After the scale-up is successful, the cluster status changes to **Scaled-up-success**. Otherwise, contact O&M personnel.
- Step 13** Start all services and set parameters on the active Master node by referring to [Step 1](#) to [Step 11](#) in the **Operations After the Master Node Specifications Scale-up** part.
- Step 14** On the **Scale Up Master Node Specifications** page, select **Are you sure you have started all services on the active Master node?** and click **OK** to complete the scale-up.

----End

#### Operations After the Master Node Specifications Scale-up

- Step 1** Log in to Manager and go to the cluster management page. For details, see [Accessing MRS Manager](#).
- Step 2** Click **Hosts**. Check basic information about the host corresponding to the business IP address of the NameNode recorded in [Step 9](#) in the **Preparing for Scaling Up Master Node Specifications** part. If the **Running Status** is **Good** and **Disk**, **Memory**, and **CPU** have values, perform [Step 9](#). If any of the preceding conditions is not met, go to the next step.
- Step 3** Log in to the standby Master node remotely. For details, see [Logging In to an MRS Cluster Node](#).
- Step 4** Run the following command to switch to user **omm**:

```
su - omm
```

**Step 5** Run the following command to start the Agent:

```
sh /opt/Bigdata/nodeagent/bin/start-agent.sh
```

**Step 6** Run the following command to check whether the Agent is started successfully:

```
jps | grep NodeAgent
```

**Step 7** Log in to Manager and go to the cluster management page. For details, see [Accessing MRS Manager](#).

**Step 8** Click **Hosts**. Check basic information about the host corresponding to the business IP address of the NameNode recorded in [Step 9](#) in the **Preparing for Scaling Up Master Node Specifications** part to ensure that **Running Status** is **Good** and **Disk, Memory**, and **CPU** have values.

 **NOTE**

It may take 3 minutes until the host status is normal after the Agent is started successfully. Please wait. If the host status is abnormal for a long time, contact O&M personnel.


**Step 9** On the Manager page, click **Hosts** and select the checkbox of the host corresponding to the business IP address of the NameNode recorded in [Step 9](#) in the **Preparing for Scaling Up Master Node Specifications** part.

**Step 10** Choose **More > Start All Instances** and wait until all instances are started.

**Step 11** Access the NameNode web UI and check the NameNode startup status.

1. On FusionInsight Manager, choose **Cluster > Services > HDFS > Dashboard**.
2. In the **HDFS Summary** column, click **NameNode** of the active or standby node that has been scaled up on the right of **NameNode Web UI**.
3. Go to the **NameNode Web UI** page, choose **Startup Progress** in the navigation bar. After ensuring that **Percent Complete** is displayed as 100%, go to the next step, as shown in [Figure 6-44](#).

**Figure 6-44** NameNode startup status



 **NOTE**

Perform [Step 11](#) for an analysis cluster and skip this step for a streaming cluster.

----End

## 6.6.6 Synchronizing Disk Information of an MRS Cluster Node

If the information displayed on the console is not the actual disk information or "Data disk: -- (Synchronize disk information)" is displayed in the node list, you can use this function to update disk information.

To obtain the latest EVS disk status, you need to synchronize disk information on the MRS console. This function updates the disk information in the cluster.

## Constraints

- Only cloud disk information can be synchronized.
- If disk information is being synchronized, the cluster cannot be scaled or upgraded.

## Synchronizing Disk Information

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS details page, click **Nodes**.

**Step 4** Click **Synchronize Disk Info**.

Wait until "Disk information synchronization request issued successfully." is displayed in the upper right corner of the page.

----End

## 6.6.7 Adding a Tag to an MRS Cluster/Node

Tags are used to identify clusters/nodes. Adding tags to clusters/nodes can help you identify and manage your resources.

- **Cluster tags:** You can add up to 20 tags to a cluster during cluster creation or add them on the details page of a created cluster. Updating a cluster tag will synchronize the tag to all nodes in the cluster.
- **Node tags:** You can use the default tag or add tags to nodes in an MRS cluster when you configure an auto scaling policy. Node tags take the tag quotas. You can view the tags of a node in the **Nodes** tab on MRS console.
- **Default tags:** An MRS cluster contains multiple nodes, and each node is an ECS and contains EVS disks. After the default tag is enabled, the system automatically creates a cluster tag and a tag for each node. The default tag is automatically synchronized to the corresponding ECS or EVS instances.

To view node tags, go to the **Nodes** tab on the MRS console, and move the cursor to the tag icon of a node in the node list.

### NOTE

- MRS tag updates are synchronized to the ECSs or EVS disks in the cluster. However, if you modify MRS cluster tags on the ECS or EVS console, the modification will not be synchronized to MRS. To ensure tag consistency, do not modify MRS cluster tags on the ECS or EVS console.
- You can add a maximum of 20 tags to a cluster. If the number of tags of a node in the cluster reaches the upper limit, no more tags can be added to the cluster.
- If default tags are enabled, a default tag is added to the cluster and each node, which takes two quotas. That is, a maximum of 20 tags can be added by default. In this case, a maximum of 18 more tags can still be added.

If your organization has configured tag strategies for MRS, add tags to clusters/nodes based on the strategies. If a tag does not comply with the tag strategies, the cluster/node may be failed to be created. Contact the organization administrator to learn more about the tag strategies.

A tag consists of a tag key and a tag value. [Table 6-30](#) provides tag key and value requirements.


**Table 6-30** Tag key and value requirements

| Parameter | Requirement                                                                                                                                                                                                                                       | Example      |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| Key       | <p>A tag key cannot be left blank.</p> <p>A tag key must be unique in a cluster.</p> <p>A tag key can contain a maximum of 128 characters.</p> <p>A tag value cannot contain special characters (=*&lt;&gt;\, /) or start or end with spaces.</p> | Organization |
| Value     | <p>A tag value can contain a maximum of 255 characters.</p> <p>A tag value cannot contain special characters (=*&lt;&gt;\, /) or start or end with spaces. This parameter can be left blank.</p>                                                  | Apache       |

## Adding Tags to a Cluster


- Adding cluster tags during cluster creation
  - a. Log in to the MRS console.
  - b. Click **Buy Cluster**. The page for buying a cluster is displayed.
  - c. Click the **Custom Config** tab.
  - d. Configure the cluster software and hardware by referring to [Manually Buying an MRS Cluster](#).
  - e. Select **Configure** on the right of **Set Advanced Options** and enter the key and value of a new tag.

**Figure 6-45** Adding a tag to a cluster




- Adding tags to an existing cluster
  - a. Log in to the MRS console.
  - b. On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
  - c. Click the **Tags** tab.
  - d. Click **Add/Edit Tag**. If this is your first time adding a tag, click **Add Tag**. In the displayed dialog box, enter the key and value of a tag, and click **Add**.

**Figure 6-46** Adding/Editing a Tag



**Add/Edit Tag** ×

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#) 

To add a tag, enter a tag key and a tag value below.

10 tags available for addition.

 **NOTE**

You can also add cluster tags by enabling default tags. All nodes will be tagged with the cluster ID and node IDs, which takes two quotas.

- e. Click **OK**.

## Adding Tags to a Node

- Node tags are automatically added when a default tag is added to a cluster. For details, see [Adding tags to an existing cluster](#).
- Adding node tags for auto scaling

If you add a tag when configuring an auto scaling policy, MRS automatically adds the tag to the new nodes and synchronizes the tag to the ECSs and EVS disks.

  - a. Log in to the MRS console.
  - b. On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
  - c. On the page that is displayed, click the **Auto Scaling** tab.



- d. Click **Edit** on the right of an existing auto scaling policy. In the displayed dialog box, enter the key and value of the tag you want to add, and click **Add**.

 **NOTE**

- You need to enable the auto scaling policy and configure scale-out rules. Otherwise, the node tags will not take effect.
  - If tag quotas are insufficient, delete the cluster tag or modify existing a tag of the auto scaling policy, and then enable the policy.
  - Tags cannot be added to auto scaling policies of resource pools.
- e. Click **OK**.

## Searching for Target Clusters by Tags

On the **Active Clusters** page, search for the target cluster by tag key or tag value.

1. Log in to the MRS console.
2. In the upper right corner of the **Active Clusters** page, click **Search by Tag** to access the search page.
3. Enter the tag of the cluster to be searched.

You can select a tag key or tag value from their drop-down lists. When the tag key or tag value is exactly matched, the system can automatically locate the target cluster. If you enter multiple tags, their intersections are used to search for the cluster.

4. Click **Search**.

The system searches for the target cluster by tag key or value.

## Managing Tags

You can view, add, and delete tags on the **Tags** tab page of the cluster.

1. Log in to the MRS console.
2. On the **Active Clusters** page, click the name of a cluster for which you want to manage tags.

The cluster details page is displayed.

3. Click the **Tags** tab and view, add, and delete tags on the tab page.

- View

On the **Tags** tab page, you can view details about tags of the cluster, including the number of tags and the key and value of each tag.

- Add

Click **Add/Edit Tag**. If this is your first time adding a tag, click **Add Tag**. In the displayed dialog box, enter the key and value of a tag, and click **OK**.

- Delete

Locate the row that contains the tag you want to delete and click **Delete** in the **Operation** column. In the displayed **Delete Tag** dialog box, enter **DELETE**, and click **Yes**.

## 6.6.8 Configuring Bootstrap Actions for an MRS Cluster Node

### 6.6.8.1 MRS Bootstrap Action Overview

You can run bootstrap actions to install additional third-party software, modify the cluster running environment, and perform other customizations.

Bootstrap actions can execute scripts on specified nodes before or after the first startup of cluster components. You can only manually run the third-party component installation script on the node to install a running cluster component.

If you choose to run bootstrap actions when scaling out a cluster, the bootstrap actions will be run on the newly added nodes in the same way. If auto scaling is enabled in a cluster, you can add an automation script in addition to configuring a resource plan. Then the automation script executes the corresponding script on the nodes that are scaled out or in to implement custom operations.

#### NOTE

- MRS 2.x and earlier versions: Bootstrap action scripts must be executed as user **root**. Otherwise, your cluster may become unavailable. You can run the **su - xxx** command to switch users in the script.
- MRS 3.x or later: Bootstrap action scripts must be executed as user **omm**. Otherwise, your cluster may become unavailable. You can run the **su - xxx** command to switch users in the script.

MRS determines the result based on the return code after the execution of the bootstrap action script. If the return code is **0**, the script is executed. If the return code is not **0**, the execution fails. If a bootstrap action script fails to be executed on a node, the corresponding boot script will fail to be executed. In this case, you can set **Action upon Failure** to choose whether to continue to execute the subsequent scripts.

- Example 1: If you set **Action upon Failure** to **Continue** for all scripts during cluster creation, all the scripts will be executed regardless of whether they are successfully executed, and the startup process will be complete.
- Example 2: If a script fails to be executed and **Action upon Failure** is set to **Stop**, subsequent scripts will not be executed and cluster creation or scale-out will fail.

You can add a maximum of 18 bootstrap actions, which will be executed before or after the cluster component is started in the order you specified. The bootstrap actions performed before or after the component startup must be completed within 60 minutes. Otherwise, the cluster creation or scale-out will fail.

### 6.6.8.2 Preparing the Bootstrap Action Script for an MRS Node

Currently, bootstrap actions support Linux shell scripts only. Script files must end with **.sh**.

Follow these steps to prepare the bootstrap action script for the MRS node:

- Step 1** Upload the required installation packages to the OBS file system.

Before compiling a script, you need to upload all required installation packages, configuration packages, and relevant files to the OBS file system in the same region.

Different regions are isolated from each other. Therefore, MRS VMs cannot download OBS files from other regions.

**Step 2** Use a script for downloading files from the OBS file system.

You can specify the file to be downloaded from OBS in the script. To upload a file to a private file system, you need to use **hadoop fs** to download the file.

To download the **myfile.tar.gz** file from the **obs://yourbucket/** directory to your local host and decompress it to the **/your-dir** directory, run the following commands:

```
source /opt/Bigdata/client/bigdata_env;hadoop fs -D fs.obs.endpoint=<obs-
endpoint> -D fs.obs.access.key=<your-ak> -D fs.obs.secret.key=<your-sk> -
copyToLocal obs://yourbucket/myfile.tar.gz ./
```

```
mkdir -p /<your-dir>
```

```
tar -zxvf myfile.tar.gz -C /<your-dir>
```

#### NOTE

- **/opt/Bigdata/client** indicates the client path. Change it based on the site requirements.
- The Hadoop client has been preinstalled on the MRS node. You can run the **hadoop fs** command to download or upload data from or to OBS.
- Obtain the obs-endpoint of each region. For details, see [Regions and Endpoints](#).
- Commands carrying authentication passwords pose security risks. Disable historical command recording before running such commands to prevent information leakage.

**Step 3** Upload the script to the OBS file system.

After script compilation, upload the script to the OBS file system in the same region. At the time you specify, each node in the cluster downloads the script from OBS and executes it as user **root**.

----End

### 6.6.8.3 Adding MRS Node Bootstrap Actions and Installing Third-Party Software

#### Prerequisites

You have prepared the bootstrap action script by referring to [Preparing the Bootstrap Action Script for an MRS Node](#).

#### Adding a Bootstrap Action When Creating a Cluster

**Step 1** Go to the [Buy Cluster](#) page.

**Step 2** Click **Custom Config**.

**Step 3** Configure the cluster software and hardware by referring to [Manually Buying an MRS Cluster](#).

**Step 4** In the **Set Advanced Options** area, select **Configure** and click **Add** in the **Bootstrap Action** area.

**Table 6-31** Parameter description

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                | Name of a bootstrap action script<br>The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and must not start with a space.<br>The value can contain 1 to 64 characters.<br><b>NOTE</b><br>A name must be unique in the same cluster. You can set the same name for different clusters.                                                                                                                        |
| Script Path         | Script path. The value can be an OBS file system path or a local VM path. <ul style="list-style-type: none"><li>• An OBS file system path must start with <b>obs://</b> and end with <b>.sh</b>, for example, <b>obs://mrs-samples/xxx.sh</b>.</li><li>• A local VM path must start with a slash (/) and end with <b>.sh</b>.</li></ul> <b>NOTE</b><br>A path must be unique in the same cluster, but can be the same for different clusters. |
| Parameter           | Bootstrap action script parameters                                                                                                                                                                                                                                                                                                                                                                                                            |
| Execution Node      | Select a type of the node where the bootstrap action script is executed.                                                                                                                                                                                                                                                                                                                                                                      |
| Executed            | Select the time when the bootstrap action script is executed. <ul style="list-style-type: none"><li>• Before initial component start</li><li>• After initial component start</li></ul> <b>NOTE</b><br>You can only manually run the third-party component installation script on the node to install a running cluster component.                                                                                                             |
| Action upon Failure | Whether to continue to execute subsequent scripts and create a cluster after the script fails to be executed.<br><b>NOTE</b><br>You are advised to set this parameter to <b>Continue</b> in the debugging phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.                                                                                                            |
| Run as root         | Whether to escalate the permission to user <b>root</b><br>If the bootstrap action requires root user operations, enable this function, or the bootstrap action may fail to execute.<br><b>NOTE</b><br>This parameter applies only to clusters of MRS 3.1.5, MRS 3.3.0, and later versions.                                                                                                                                                    |

**Step 5** Click **OK**.

After the bootstrap action is added, you can edit, clone, or delete it in the **Operation** column.

----End

## Adding a Bootstrap Action to an Existing Cluster

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the page that is displayed, click the **Bootstrap Actions** tab.

**Step 4** Click **Add** and set parameters as prompted.

**Figure 6-47** Add Bootstrap Action

**Add Bootstrap Action** ×

\* Name

\* Script Path

\* Execution Node  Master  Analysis Core  
 Streaming Core  Analysis Task  
 Streaming Task  
Active Master

Parameter  0/128  
0/128

\* Executed  ▼

\* Action upon Failure  Continue  Stop

**Table 6-32** Parameter description

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                | <p>Name of a bootstrap action script</p> <p>The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and must not start with a space.</p> <p>The value can contain 1 to 64 characters.</p> <p><b>NOTE</b><br/>A name must be unique in the same cluster. You can set the same name for different clusters.</p>                                                                                                                      |
| Script Path         | <p>Script path. The value can be an OBS file system path or a local VM path.</p> <ul style="list-style-type: none"> <li>• An OBS file system path must start with <b>obs://</b> and end with <b>.sh</b>, for example, <b>obs://mrs-samples/xxx.sh</b>.</li> <li>• A local VM path must start with a slash (/) and end with <b>.sh</b>.</li> </ul> <p><b>NOTE</b><br/>A path must be unique in the same cluster, but can be the same for different clusters.</p> |
| Parameter           | Bootstrap action script parameters                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Execution Node      | Select a type of the node where the bootstrap action script is executed.                                                                                                                                                                                                                                                                                                                                                                                        |
| Executed            | <p>Select the time when the bootstrap action script is executed.</p> <ul style="list-style-type: none"> <li>• Before initial component start</li> <li>• After initial component start</li> </ul> <p><b>NOTE</b><br/>You can only manually run the third-party component installation script on the node to install a running cluster component.</p>                                                                                                             |
| Action upon Failure | <p>Whether to continue to execute subsequent scripts and create a cluster after the script fails to be executed.</p> <p><b>NOTE</b><br/>You are advised to set this parameter to <b>Continue</b> in the debugging phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.</p>                                                                                                                  |
| Run as root         | <p>Whether to escalate the permission to user <b>root</b></p> <p>If the bootstrap action requires root user operations, enable this function, or the bootstrap action may fail to execute.</p> <p><b>NOTE</b><br/>This parameter applies only to clusters of MRS 3.1.5, MRS 3.3.0, and later versions.</p>                                                                                                                                                      |

**Step 5** Click **OK** to save the configuration.

**Step 6** Click **Yes**. After a bootstrap action is added, you can modify or delete it in the bootstrap action list.

----End

## Adding an Automatic Script for Auto Scaling Nodes

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** Click **Auto Scaling**.

**Step 4** Click **Configure Automation Script**.

**Step 5** Click **Add**. Set parameters by referring to [Table 6-32](#).

**Step 6** Click **OK** to save the automation script configurations.

----End

### 6.6.8.4 Viewing the Bootstrap Action Execution Records of an MRS Node

You can view the execution result of the bootstrap operation on the **Bootstrap Action** page.

## Viewing the Execution Result

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the cluster details page, click **Bootstrap Action**. Information about the bootstrap actions added during cluster creation is displayed.

**Figure 6-48** Bootstrap action information

| Name                                 | Script Path                | Execution Node               | Parameter | Execution Time                | Action upon Fail. | Executed | Status  |
|--------------------------------------|----------------------------|------------------------------|-----------|-------------------------------|-------------------|----------|---------|
| <input type="checkbox"/> bootstrap_0 | s3a://bucket/program/vx.sh | Active Master/Standby Master | --        | After initial component start | Continue          | --       | PENDING |
| <input type="checkbox"/> 232-2-w2    | s3a://bucket/program/vx.sh | Active Master/Standby Master | --        | After initial component start | Continue          | --       | PENDING |

### NOTE

- You select **Before initial component start** or **After initial component start** in the upper right corner to query information about the related bootstrap actions.
- The last execution result is listed here. For a newly created cluster, the records of bootstrap actions executed during cluster creation are listed. If a cluster is expanded, the records of bootstrap actions executed on the newly added nodes are listed.

----End

## Viewing Execution Logs

If you want to view the run logs of a bootstrap action, set **Action upon Failure** to **Continue** when adding the bootstrap action. And then, log in to each node to view the run logs in the `/var/log/Bootstrap` directory.

If you add bootstrap actions before and after component start, you can distinguish bootstrap action logs of the two phases based on the timestamps.

You are advised to print logs in detail in the script so that you can view the detailed run result. MRS redirects the standard output and error output of the script to the log directory of the bootstrap action.

## 6.7 Managing the MRS Cluster Client

### 6.7.1 Updating the MRS Cluster Client After the Server Configuration Expires

An MRS cluster provides a client for you to connect to a server, view task results, or manage data. If you modify service configuration parameters on Manager and restart the service, you need to download and install the client again or use the configuration file to update the client.

For clusters of MRS 2.x or earlier, the original client is stored in the `/opt/client` directory of all nodes in the cluster by default when you create a cluster. After the cluster is created, only the client of a Master node can be directly used. To use the client of a Core node, you need to update the client configuration file first.

### Updating Client Configurations (Version 3.x or Later)

#### Method 1

**Step 1** Log in to MRS Manager by referring to [Accessing MRS Manager](#). Choose **Cluster > Overview > More > Download Client**.

**Step 2** Choose **More > Download Client > Configuration Files Only**.

The generated compressed file contains the configuration files of all services.



**Figure 6-49** Downloading the client configuration file

## Download Cluster Client

Download the **MRS Cluster** client. The cluster client provides all services.

Select Client Type:  Complete Client  Configuration Files Only

Select Platform Type:  x86\_64  aarch64

Save to Path :  ?

- Step 3** Determine whether to generate a configuration file on the cluster node.
- If yes, select **Save to Path**, and click **OK** to generate the client file. By default, the client file is generated in **/tmp/FusionInsight-Client** on the active management node. You can also store the client file in other directories, and user **omm** has the read, write, and execute permissions on the directories. Then go to **Step 4**.
  - If no, click **OK**, specify a local save path, and download the complete client. Wait until the download is complete and go to **Step 4**.
- Step 4** Use WinSCP to save the compressed file to the client installation directory, for example, **/opt/hadoopclient**, as the client installation user.
- Step 5** Decompress the software package.
- Run the following commands to go to the directory where the client is installed, and decompress the file to a local directory. For example, the downloaded client file is **FusionInsight\_Cluster\_1\_Services\_Client.tar**.
- ```
cd /opt/hadoopclient
tar -xvf FusionInsight_Cluster_1_Services_Client.tar
```
- Step 6** Verify the software package.
- Run the following command to verify the decompressed file and check whether the command output is consistent with the information in the **sha256** file.
- ```
sha256sum -c
FusionInsight_Cluster_1_Services_ClientConfig_ConfigFiles.tar.sha256
```
- FusionInsight\_Cluster\_1\_Services\_ClientConfig\_ConfigFiles.tar: OK
- Step 7** Decompress the package to obtain the configuration file.
- ```
tar -xvf FusionInsight_Cluster_1_Services_ClientConfig_ConfigFiles.tar
```
- Step 8** Run the following command in the client installation directory to update the client using the configuration file:
- ```
sh refreshConfig.sh Client installation directory Directory where the configuration file is located
```

For example, run the following command:

```
sh refreshConfig.sh /opt/hadoopclient /opt/hadoopclient/
FusionInsight_Cluster_1_Services_ClientConfig_ConfigFiles
```

If the following information is displayed, the configurations have been updated successfully.

```
Succeed to refresh components client config.
```

----End

#### Method 2:

- Step 1** Log in to the client installation node as user **root**.
- Step 2** Go to the client installation directory, for example, **/opt/hadoopclient** and run the following commands to update the configuration file:

```
cd /opt/hadoopclient
sh autoRefreshConfig.sh
```

- Step 3** Enter the username and password of the FusionInsight Manager administrator and the floating IP address of OMS.

#### NOTE

To obtain the floating IP address of OMS, log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of OMS. Record the value of **inet**. If the floating IP address of OMS cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one Master node, query and record the IP address on the Master node.

- Step 4** Enter the names of the components whose configuration needs to be updated. Use commas (,) to separate the component names. Press **Enter** to update the configurations of all components if necessary.

If the following information is displayed, the configurations have been updated successfully.

```
Succeed to refresh components client config.
```

----End

## Updating Client Configurations (Version 2.x or Earlier)

### Method 1 (applicable to all versions)

- Step 1** Log in to MRS Manager and choose **Services**.
- Step 2** Click **Download Client**.

Set **Client Type** to **Only configuration files**, **Download To** to **Server**, and click **OK** to generate the client configuration file. The generated file is saved in the **/tmp/MRS-client** directory on the active management node by default. You can customize the file path.

**Figure 6-50** Downloading the client configuration file

## Download Client

Warning: Generating a client will occupy a large number of disk I/Os. You are advised not to download a client when the cluster is being installed, started, and patched, or in other unstable states.

\* Client Type  All client files  Only configuration files

\* Download To  Server  Remote host

Files will only be saved to the following path on the server. Existing client files in the path will be overwritten.

\* Client Path

OK

Cancel

**Step 3** Query and log in to the active Master node.

**Step 4** If you use the client in the cluster, run the following command to switch to user **omm**. If you use the client outside the cluster, switch to user **root**.

```
sudo su - omm
```

**Step 5** Go to the client directory.

```
cd Client installation directory
```

**Step 6** Update client configurations.

```
sh refreshConfig.sh Client installation directoryFull path of the client configuration file package
```

Example:

```
sh refreshConfig.sh /opt/Bigdata/client /tmp/MRS-client/
MRS_Services_Client.tar
```

If the following information is displayed, the configurations have been updated successfully.

```
ReFresh components client config is complete.
Succeed to refresh components client config.
```

----End

### Method 2:

**Step 1** After the cluster is installed, run the following command to switch to user **omm**. If you use the client outside the cluster, switch to user **root**.

```
sudo su - omm
```

**Step 2** Go to the client directory.

```
cd Client installation directory
```

**Step 3** Run the following command and enter the name of an MRS Manager user with the download permission and its password (for example, the username is **admin** and the password is the one set during cluster creation) as prompted to update client configurations.

```
sh autoRefreshConfig.sh
```

**Step 4** After the command is executed, the following information is displayed, where *XXX* indicates the name of the component installed in the cluster. To update client configurations of all components, press **Enter**. To update client configurations of some components, enter the component names and separate them with commas (,).

Components "xxx" have been installed in the cluster. Please input the comma-separated names of the components for which you want to update client configurations. If you press Enter without inputting any component name, the client configurations of all components will be updated:

If the following information is displayed, the configurations have been updated:

```
Succeed to refresh components client config.
```

If the following information is displayed, the username or password is incorrect.

```
login manager failed,Incorrect username or password.
```

#### NOTE

- This script automatically connects to the cluster and invokes the **refreshConfig.sh** script to download and update the client configuration file.
- By default, the client uses the floating IP address specified by **wsom=xxx** in the **Version** file in the installation directory to update the client configurations. To update the configuration file of another cluster, modify the value of **wsom=xxx** in the **Version** file to the floating IP address of the corresponding cluster before performing this step.

----End

## 6.7.2 Viewing the Installed MRS Cluster Client

FusionInsight Manager supports unified management of cluster client installation information. After a user downloads and installs a client, FusionInsight Manager automatically records information about the installed (registered) client to facilitate query and management.

In addition, you can manually add or modify the information about clients that are not automatically registered, for example, clients installed in earlier versions.

#### NOTE

This section applies only to MRS 3.x or later.

### Viewing the Installed Cluster Client

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Cluster** and choose **Client Management** to view information about clients installed in the cluster.

You can view the IP address, installation path, component list, registration time, and installation user of the node where the client is located.

When the client is downloaded and installed in the cluster of the latest version, the client information is automatically registered.

**Figure 6-51** Client information

Client Management



The screenshot shows the 'Client Management' page with a table of registered clients. The table has columns for IP Address, Installation Path, Component, Registration Date, User, Platform Type, Version, Registration Method, and Operation. A single client is listed with the path /opt/client, registered on Sep 29, 2021, by user root, on an x86\_64 platform, with an automatic registration method. The interface includes buttons for 'Add', 'Export All', and 'Batch Delete', along with an 'Advanced Search' dropdown and refresh/refresh icons.

| IP Address | Installation ... | Component... | Registration...  | User | Platform Type | Version | Registration M... | Operation     |
|------------|------------------|--------------|------------------|------|---------------|---------|-------------------|---------------|
|            | /opt/client      |              | Sep 29, 2021 ... | root | x86_64        |         | Automatic         | Delete Modify |

- Step 3** To manually add information about an installed client, click **Add** and manually add the IP address, installation path, user, platform information, and registration information of the client as prompted.

Configure the client information and click **OK**.

You can also manually modify or delete the information about the manually registered client.

- Step 4** On the **Client Management** page, click **Export All** to export information about all registered clients to the local PC.

 **NOTE**

On the **Client Management** page, only components that have clients are displayed in the component list. Therefore, some components that do not have clients and have special components are not displayed.

----End

## 6.7.3 Batch Upgrading MRS Cluster Clients

The client package downloaded from FusionInsight Manager contains the client batch upgrade tool. When multiple clients need to be upgraded after the cluster upgrade or scale-out, you can use this tool to upgrade the clients in batches with a few clicks. In addition to upgrading clients in batches, this tool also provides the lightweight function of updating the `/etc/hosts` files on the nodes where the clients are located in batches.

 **NOTE**

This section applies only to MRS 3.x or later.

### Preparing for the Batch Upgrade

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster > Overview > More > Download Client** (for MRS 3.3.0 or later, click **Download Client** on the **Homepage** page) to download the client to a specified directory on the server.

For details, see [Installing an MRS Cluster Client](#).

Decompress the downloaded client package and find the **batch\_upgrade** directory, for example, `/tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_ClientConfig/batch_upgrade`.

- Step 3** Choose **Cluster > Client Management**. On the **Client Management** page, click **Export All** to export all client information to the local PC.
- Step 4** Decompress the exported client information and upload the **client-info.cfg** file to the **batch\_upgrade** directory.
- Step 5** Perform the following operations to supplement the missing passwords in the **client-info.cfg** file:

Run the **vi client-info.cfg** command to add a user password.

Example:

```
clientIp,clientPath,user,password
10.10.10.100,/home/omm/client /home/omm/client2,omm,Password
```

The fields in the configuration file are as follows:

- **clientIp**: indicates the IP address of the node where the client is located.
- **clientPath**: indicates the client installation path. Multiple paths are separated by spaces. Note that the path cannot end with a slash (/).
- **user**: indicates the username of the node.
- **password**: indicates the user password of the node.

 **NOTE**

- If the execution fails, view the **node.log** file in the **work\_space/log\_XXX** directory.
- There can be security risks if a configuration file contains the authentication password. You are advised to delete the configuration file or use other secure methods to keep the password.

----End

## Procedure

- Step 1** Log in to the client download node as the user who wants to install the client.
- Step 2** Perform the upgrade.

```
sh client_batch_upgrade.sh -u -f /tmp/FusionInsight-Client/
FusionInsight_Cluster_1_Services_Client.tar -g /tmp/FusionInsight-Client/
FusionInsight_Cluster_1_Services_ClientConfig/batch_upgrade/client-info.cfg
```

---

**NOTICE**

You are advised to delete the **client-info.cfg** file as soon as possible after the upgrade because the password has been configured.

---

- Step 3** After the upgrade is complete, verify the upgrade result by running the **sh client\_batch\_upgrade.sh -c** command.
- Step 4** If the client is faulty, run the **sh client\_batch\_upgrade.sh -s** command to roll back the client.

 NOTE

- The client batch upgrade tool moves the original client to the backup directory, and then uses the client package specified by the **-f** parameter to install the client. Therefore, if the original client contains customized content, manually save the customized content from the backup directory or move the customized content to the client directory after the upgrade before running the **-c** command. The backup path on the client is *{Original client path}-backup*.
- The **-u** command is the prerequisite for the **-c** and **-s** commands. You can run the **-c** command to commit the upgrade or the **-s** command to perform a rollback only after the **-u** command is executed to perform an upgrade.
- You can run the **-u** command multiple times to upgrade only the clients that fail to be upgraded.
- The client batch upgrade tool also supports the clients of earlier versions.
- When upgrading a client installed by a non-root user, ensure that the user has the read and write permissions on the directory where the client is located and the parent directory on the target node. Otherwise, the upgrade will fail.
- The client package specified by the **-f** parameter must be a full client package. The client packages of a single component or some components cannot be used as the input.

----End

## Updating the hosts Files in Batches

**Step 1** Complete all preparations by referring to [Preparing for the Batch Upgrade](#).

**Step 2** Check whether the user configured for the node where the **/etc/hosts** files need to be updated is **root**.

- If yes, go to [Step 3](#).
- If no, change the user to **root** and go to [Step 3](#).

**Step 3** Update the **/etc/hosts** files on the nodes where the client is located in batches.

```
sh client_batch_upgrade.sh -r -f /tmp/FusionInsight-Client/
FusionInsight_Cluster_1_Services_Client.tar -g /tmp/FusionInsight-Client/
FusionInsight_Cluster_1_Services_ClientConfig/batch_upgrade/client-info.cfg
```

 NOTE

- When you batch update the **/etc/hosts** files, the entered client package can be a complete client package or a client package that contains only configuration files (recommended).
- The user configured for the host where the **/etc/hosts** files need to be updated must be **root**. Otherwise, the update fails.

----End

## 6.8 Managing MRS Cluster Jobs

### 6.8.1 Stopping and Deleting an MRS Cluster Job

You can manually stop a running MRS job on the console.

After a job is executed, you can delete it if you do not need to view the job information.

## Stopping an MRS Cluster Job

Spark SQL jobs cannot be stopped. After a job is stopped, its status changes to **Terminated** and the job cannot be executed again.

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** Click **Jobs**.
- Step 4** Select a running job, and choose **More > Stop** in the **Operation** column.  
The job status changes from **Running** to **Terminated**.  
----End

## Deleting an MRS Cluster Job

You can delete a single job or multiple jobs in batches. A deleted job cannot be restored. Therefore, exercise caution when deleting a job.

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** Click **Jobs**.
- Step 4** Choose **More > Delete** from the **Operation** in the row of the target job to be deleted. Enter **DELETE** in the **Delete Job** dialog box and click **OK**.  
This step deletes only a single job.
- Step 5** Select multiple jobs and click **Delete** on the upper left of the job list. Enter **DELETE** in the **Delete Job** dialog box and click **OK**.  
You can delete one, multiple, or all jobs.  
----End

## 6.8.2 Configuring Notification Rules for MRS Jobs

MRS uses SMN to offer a publish/subscribe model to achieve one-to-multiple message subscriptions and notifications in a variety of message types (SMSs and emails). You can configure job notification rules to receive notifications immediately upon a job execution success or failure.


- Step 1** Log in to the management console.
- Step 2** Click **Service List**. Under **Management & Governance**, click **Simple Message Notification**.
- Step 3** Create a topic and add subscriptions to the topic. For details, see [Configuring Notifications for MRS Cluster Alarms and Events](#).



**Step 4** Go to the MRS management console, and click the cluster name to go to the cluster details page.

**Step 5** Click the **Alarms** tab, and choose **Notification Rules > Add Notification Rule**.

**Figure 6-52** Adding a notification rule



**Step 6** Configure a notification rule for sending job execution results to subscribers.

**Table 6-33** Parameters of adding a notification rule

| Parameter            | Description                                                                                                                                                                                             |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rule Name            | User-defined notification rule name. Only digits, letters, hyphens (-), and underscores (_) are allowed.                                                                                                |
| Message Notification | If you enable this function, subscription messages will be sent to subscribers.                                                                                                                         |
| Topic Name           | Select an existing topic or click <b>Create Topic</b> to create a topic.                                                                                                                                |
| Notification Type    | Select <b>Event</b> .                                                                                                                                                                                   |
| Subscription Items   | <ol style="list-style-type: none"> <li>Click  next to <b>Suggestion</b>.</li> <li>Click  next to <b>Manager</b>.</li> <li>Select <b>Job Running Succeeded</b> and <b>Job Running Failed</b>.</li> </ol> |

----End

## 6.9 Managing MRS Cluster Tenants

### 6.9.1 Introduction to MRS Multi-Tenancy

#### Overview

- Background:**

Modern enterprises' data clusters are becoming more and more centralized and cloud-based. Enterprise-class big data clusters must meet the following requirements:

  - Carry data of different types and formats and run jobs and applications of different types (such analysis, query, and stream processing).

- Isolate data of a user from that of another user who has demanding requirements on data security, such as a bank or government institute.

The preceding requirements bring the following challenges to the big data clusters:

- Proper allocation and scheduling of resources to ensure stable operating of applications and jobs.
- Strict access control to ensure data and service security.

Multi-tenancy isolates the resources of a big data cluster into resource sets. Users can lease desired resource sets to run applications and jobs and store data. In a big data cluster, multiple resource sets can be deployed to meet diverse requirements of multiple users.

MRS provides a complete enterprise-class big data multi-tenant solution.

- **Multi-tenancy:**

An MRS cluster provides various resources and services that can be shared among organizations, departments, or applications. The cluster provides logical entities, that is, tenants, to use these resources and services. Currently, only the analysis cluster supports tenant management.

A mode involving different tenants is called multi-tenant mode. Multi-tenancy refers to multiple resource sets (a resource set is a tenant) in the MRS cluster and is able to allocate and schedule resources. The resources include computing resources and storage resources. The MRS cluster offers multi-tenancy, supporting a layered tenant model and enabling dynamic tenant creation or deletion to isolate resources. It dynamically allocates and configures compute and storage resources for each tenant.

The computing resources indicate tenants' Yarn task queue resources. The task queue quota can be modified, and the task queue usage status and statistics can be viewed.

The storage resources can be stored on HDFS. You can add and delete the HDFS storage directories of tenants, and set the quotas of file quantity and the storage space of the directories.

Tenants can create and manage tenants in a cluster based on service requirements.

- Roles, computing resources, and storage resources are automatically created when tenants are created. By default, all permissions of the new computing resources and storage resources are allocated to a tenant's roles.
- Permissions to view the current tenant's resources, add a subtenant, and manage the subtenant's resources are granted to the tenant's roles by default.
- After you have modified the tenant's computing or storage resources, permissions of the tenant's roles are automatically updated.

MRS supports a maximum of 512 tenants. The default tenants created by the system include **default**. Tenants that are in the topmost layer with the default tenant are called level-1 tenants.

- **Resource pool:**

YARN task queues support only one scheduling policy, namely label based scheduling. This policy allows YARN task queues to select NodeManagers with specific node labels, ensuring that tasks run on designated nodes and utilize

target hardware resources. For example, YARN tasks requiring a large memory capacity can run on nodes labeled as large memory resources, preventing poor service performance.

In an MRS cluster, tenants logically partition nodes to form a resource pool with multiple NodeManagers. YARN task queues can be associated with specified resource pools by configuring queue capacity policies, ensuring efficient and independent resource utilization in the resource pools.

MRS supports a maximum of 50 resource pools. The system has a **default** resource pool.

- **Advantages:**

- Proper resource configuration and isolation

The resources of a tenant are isolated from those of another tenant. The resource use of a tenant does not affect other tenants. This mechanism ensures that each tenant can configure resources based on service requirements, improving resource utilization.

- Resource consumption measurement and statistics

Tenants are system resource applicants and consumers. System resources are planned and allocated based on tenants. Resource consumption by tenants can be measured and collected.

- Assured data security and access security

In multi-tenant scenarios, the data of each tenant is stored separately to ensure data security. The access to tenants' resources is controlled to ensure access security.

## Multi-Tenant Model

- **Multi-tenant model**

The following figure shows a multi-tenant model.

Figure 6-53 Multi-tenant model

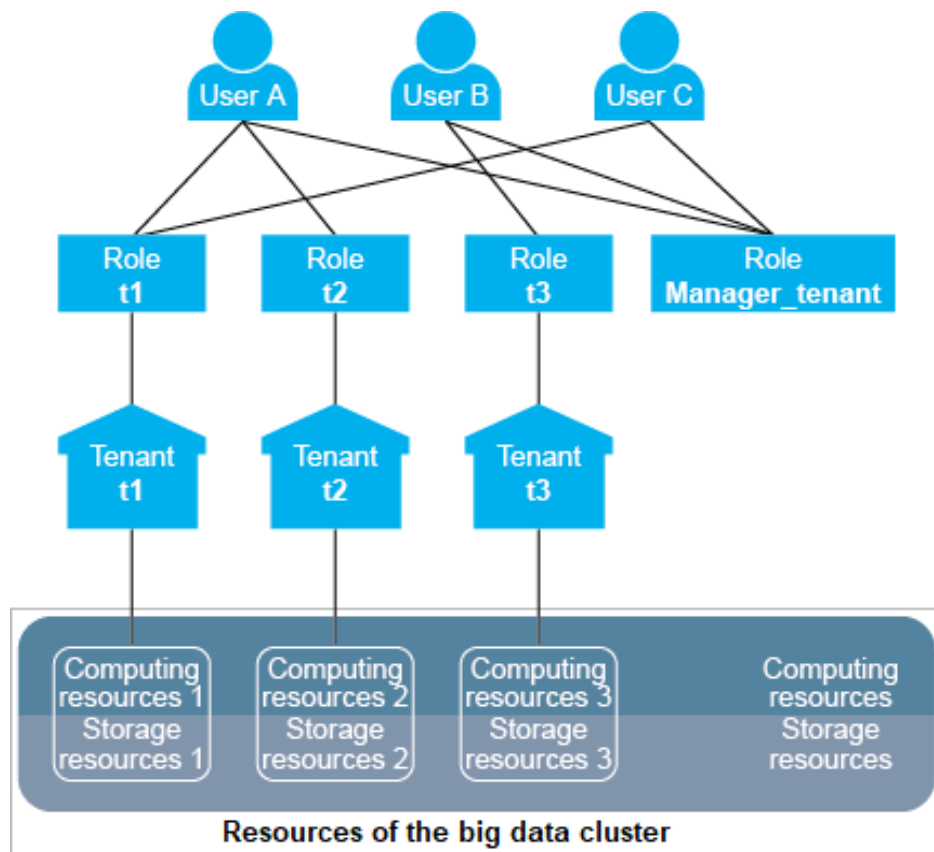


Table 6-34 describes the concepts involved in Figure 6-53.

Table 6-34 Concepts in the model

| Concept | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Users   | A natural person who has a username and password and uses the big data cluster.<br>There are three different users in the figure: users A, B, and C.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Role    | A role is a carrier of one or more permissions. Permissions are assigned to specific objects, for example, access permissions for the <code>/tenant</code> directory in HDFS.<br>There are four roles: <b>t1</b> , <b>t2</b> , <b>t3</b> , and <b>Manager_tenant</b> . <ul style="list-style-type: none"> <li>Roles <b>t1</b>, <b>t2</b>, and <b>t3</b> are automatically generated when tenants are created. The role names are the same as the tenant names. That is, roles <b>t1</b>, <b>t2</b>, and <b>t3</b> map to tenants <b>t1</b>, <b>t2</b>, and <b>t3</b>. Role names and tenant names need to be used in pair.</li> <li>Role <b>Manager_tenant</b> is defaulted in the cluster and cannot be used separately.</li> </ul> |

| Concept  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tenant   | <p>A tenant is a resource set in a big data cluster. Multiple tenants are referred to as multi-tenancy. The resource sets further divided under a tenant are called sub-tenants.</p> <p>There are three tenants: <b>t1</b>, <b>t2</b>, and <b>t3</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Resource | <ul style="list-style-type: none"> <li>• Compute resources include CPUs and memory. Compute resources are allocated to a tenant from the cluster's total resources. Each tenant's resources are isolated from others. According to the figure, compute resources 1, 2, and 3 are allocated to tenants <b>t1</b>, <b>t2</b>, and <b>t3</b> respectively from the cluster's total compute resources.</li> <li>• Storage resources include disks and third-party storage systems. Storage resources are allocated to each tenant from the cluster's total storage. Each tenant's storage is isolated from others. According to the figure, storage resources 1, 2, and 3 are allocated for tenants <b>t1</b>, <b>t2</b>, and <b>t3</b> respectively from the cluster's total storage resources.</li> </ul> |

To use a tenant's resources or add/delete sub-tenants, a user must be assigned both the tenant role and the **Manager\_tenant** role. [Table 6-35](#) lists the roles assigned to each user in [Figure 6-53](#).

**Table 6-35** Roles assigned to each user

| User   | Role                                                                                                                               | Permission                                                                                                                                                                            |
|--------|------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| User A | <ul style="list-style-type: none"> <li>• Role <b>t1</b></li> <li>• Role <b>t2</b></li> <li>• Role <b>Manager_tenant</b></li> </ul> | <ul style="list-style-type: none"> <li>• Uses the resources of tenants <b>t1</b> and <b>t2</b>.</li> <li>• Adds or deletes sub-tenants of tenants <b>t1</b> and <b>t2</b>.</li> </ul> |
| User B | <ul style="list-style-type: none"> <li>• Role <b>t3</b></li> <li>• Role <b>Manager_tenant</b></li> </ul>                           | <ul style="list-style-type: none"> <li>• Uses the resources of tenant <b>t3</b>.</li> <li>• Adds or deletes sub-tenants of tenant <b>t3</b>.</li> </ul>                               |
| User C | <ul style="list-style-type: none"> <li>• Role <b>t1</b></li> <li>• Role <b>Manager_tenant</b></li> </ul>                           | <ul style="list-style-type: none"> <li>• Uses the resources of tenant <b>t1</b>.</li> <li>• Adds or deletes sub-tenants of tenant <b>t1</b>.</li> </ul>                               |

A user can be assigned multiple roles, and one role can also be assigned to multiple users. Users are associated to tenants after being assigned the tenant roles. Therefore, tenants and users form a many-to-many relationship. Users can access resources across multiple tenants, and multiple users can share resources within the same tenant. For example, in [Figure 6-53](#), user A uses

the resources of tenants **t1** and **t2**, and users A and C uses the resources of tenant **t1**.

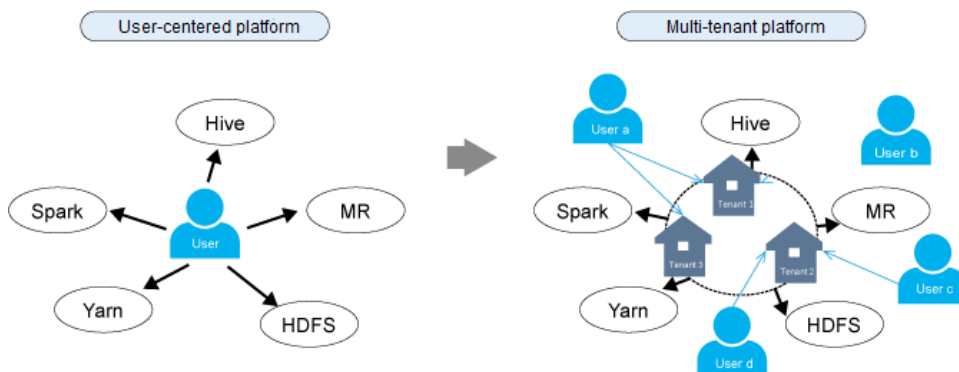
**NOTE**

The concepts of a parent tenant, sub-tenant, level-1 tenant, and level-2 tenant are all designed for the multi-tenant service scenarios. Pay attention to the differences between these concepts and the concepts of a leaf tenant resource and non-leaf tenant resource on FusionInsight Manager.

- Level-1 tenant: determined based on the tenant's level. For example, the first created tenant is a level-1 tenant and its sub-tenant is a level-2 tenant.
  - Parent tenant and sub-tenant: indicates the hierarchical relationship between tenants.
  - Non-leaf tenant resource: indicates the tenant type selected during tenant creation. This tenant type can be used to create sub-tenants.
  - Leaf tenant resource: indicates the tenant type selected during tenant creation. This tenant type cannot be used to create sub-tenants.
- **Multi-tenant platform:**

The MRS big data platform's core concept is the tenant, which enables the transition from user-centered to multi-tenant architecture to support enterprise multi-tenant applications. **Figure 6-54** shows the transformation of big data platforms.

**Figure 6-54** Platform transformation from user-centered to multi-tenant



On a user-centered big data platform, users can directly access and use all resources and services.

- However, user applications may use only partial cluster resources, resulting in low resource utilization.
- The data of different users may be stored together, decreasing data security.

On a multi-tenant big data platform, users use required resources and services by accessing the tenants.

- Resources are allocated and scheduled based on application requirements and used based on tenants, increasing resource utilization.
- Users can access the resources of tenants only after being associated with tenant roles, enhancing access security.
- The data of tenants is isolated, ensuring data security.

## Multi-Tenant Resource

MRS cluster resources are classified into compute resources and storage resources. The multi-tenant architecture implements resource isolation.

- **Compute Resources**

Compute resources include CPUs and memory. One tenant cannot occupy the compute resources of another tenant.

Compute resources are divided into static service resources and dynamic resources.

 **NOTE**

The resources allocated to YARN in a big data cluster are static service resources but can be dynamically allocated to job queues by YARN.

- **Static Service Resources**

Static service resources are compute resources allocated to each service and are not shared between services. The total compute resources of each service are fixed. These services include Flume, HBase, HDFS, and Yarn.

- **Dynamic Resources**

YARN provides distributed resource management for a big data cluster. The total volume of resources allocated to YARN can be configured. Then YARN allocates and schedules compute resources for job queues. For MapReduce, Spark, Flink, and Hive task queues, compute resources are allocated and scheduled by YARN.

YARN queues are fundamental units of scheduling compute resources.

The resources obtained by tenants using YARN queues are dynamic resources. Users can dynamically create and modify the queue quotas and view the status and statistics of the queues.

- **Resource pool:**

Enterprise IT systems often face complex cluster environments and diverse upper-layer requirements, such as the following scenarios:

- Heterogeneous cluster: The computing speed, storage capacity, and network performance of each node in the cluster are different. All the tasks of complex applications need to be properly allocated to each compute node in the cluster based on service requirements.
- Computing isolation: Data must be shared among multiple departments but computing resources must be distributed onto different compute nodes.

These require that the compute nodes be further partitioned.

Resource pools are used to specify the configuration of dynamic resources. YARN queues are associated with resource pools for resource allocation and scheduling.

One tenant can have only one default resource pool. Users can be assigned the role of a tenant to use the resources in the resource pool of the tenant. To use resources in multiple resource pools, a user can be assigned the roles of multiple tenants.

- **Dynamic resource scheduling:**

YARN dynamic resources support label-based scheduling. This policy creates labels for compute nodes (YARN NodeManagers) and adds the compute nodes with the same label into the same resource pool. Then YARN dynamically associates the queues with resource pools based on the resource requirements of the queues.

For example, a cluster has more than 40 nodes which are labeled by **Normal**, **HighCPU**, **HighMEM**, or **HighIO** based on their hardware and network configurations and added into four resource pools, respectively. [Table 6-36](#) describes the performance of each node in the resource pool.

**Table 6-36** Performance of each node in a resource pool

| Label   | Node Count | Hardware and Network Configuration | Added To        | Associated With           |
|---------|------------|------------------------------------|-----------------|---------------------------|
| Normal  | 10         | Minor                              | Resource pool A | Common queue              |
| HighCPU | 10         | High-performance CPU               | Resource pool B | Computing-intensive queue |
| HighMEM | 10         | Large memory                       | Resource pool C | Memory-intensive queue    |
| HighIO  | 10         | High-performance network           | Resource pool D | I/O-intensive queue       |

A queue can use only the compute nodes in its associated resource pool.

- A common queue is associated with resource pool A and uses **Normal** nodes with general hardware and network configurations.
- A computing-intensive queue is associated with resource pool B and uses **HighCPU** nodes with high-performance CPUs.
- A memory-intensive queue is associated with resource pool C and uses **HighMEM** nodes with large memory.
- An I/O-intensive queue is associated with resource pool C and uses **HighIO** nodes with high-performance network.

YARN queues are associated with specified resource pools to efficiently utilize resources in resource pools and maximize node performance.



FusionInsight Manager supports a maximum of 50 resource pools.  
The system has a default resource pool.

- **Storage resources**

Storage resources include disks and third-party storage systems. One tenant cannot access the data of another tenant.

As a distributed file storage service in a big data cluster, HDFS stores all the user data of the upper-layer applications in the big data cluster, including the data written to HBase tables or Hive tables.

A directory is the basic unit of allocating HDFS storage resources. HDFS supports the conventional hierarchical file structure. Users or applications can create directories and create, delete, move, or rename files in directories. Tenants can obtain storage resources from specified directories in the HDFS file system.

The storage resource scheduling mechanism is as follows:

- HDFS directories can be stored on nodes with specified labels or disks of specified hardware types.
  - When both real-time query and data analysis tasks are running in the same cluster, the real-time query tasks need to be deployed only on certain nodes, and the task data must also be stored on these nodes.
  - Based on actual service requirements, key data needs to be stored on highly reliable nodes.
- Administrators can flexibly configure HDFS data storage policies based on actual service requirements and data features to store data on specified nodes.
- For tenants, storage resources refer to the HDFS resources they use. Data of specified directories can be stored to the tenant-specified storage paths, thereby implementing storage resource scheduling and ensuring data isolation between tenants.
- Users can add or delete HDFS storage directories of tenants and set the file quantity quota and storage capacity quota of directories to manage storage resources.

## Schedulers

Multi-tenant schedulers are classified into the open-source Capacity scheduler and enhanced Superior scheduler. By default, the Superior scheduler is enabled for the MRS cluster.

- The Capacity scheduler is an open-source scheduler.
- The Superior scheduler is an enhanced version and named after the Lake Superior, indicating that the scheduler can manage a large amount of data.

### NOTE


You can check the scheduler type of YARN from the `yarn.resourcemanager.scheduler.class` value. For details about how to switch the scheduler, see [Switching the MRS Tenant Resource Scheduler](#).

To meet enterprise requirements and tackle scheduling challenges faced by the YARN community, the Superior scheduler makes the following enhancements in addition to inheriting the advantages of the Capacity scheduler and Fair scheduler:

- **Enhanced resource sharing policy**  
The Superior scheduler supports queue hierarchy. It integrates the functions of open-source schedulers and shares resources based on configurable policies. In terms of instances, administrators can use the Superior scheduler to configure an absolute value or percentage policy for queue resources. The resource sharing policy of the Superior scheduler enhances label-based scheduling of YARN as a resource pool feature. The nodes in the YARN cluster can be grouped based on the capacity or service type to ensure that queues can more efficiently utilize resources.
- **Tenant-based resource reservation policy**  
Some tenants may run critical tasks at some time, and their resource requirements must be preferentially addressed. The Superior scheduler builds a mechanism to support the resource reservation policy. Reserved resources can be allocated to the critical tasks running in the specified tenant queues in a timely manner to ensure proper task execution.
- **Fair sharing among tenants and resource pool users**  
The Superior scheduler allows shared resources to be configured for users in a queue. Each tenant may have users with different weights. Heavily weighted users may require more shared resources.
- **Ensured scheduling performance in a big cluster**  
The Superior scheduler receives heartbeats from each NodeManager and saves resource information in memory, which enables the scheduler to control cluster resource usage globally. The Superior scheduler uses the push scheduling model, which makes the scheduling more precise and efficient and remarkably improves cluster resource utilization. Additionally, the Superior scheduler delivers excellent performance when the interval between NodeManager heartbeats is long and prevents heartbeat storms in big clusters.
- **Priority policy**  
If the minimum resource requirement of a service cannot be met after the service obtains all available resources, a preemption occurs. The preemption function is disabled by default.

## Multi-Tenant Management

- **Unified multi-tenant management**  
Log in to the MRS console or FusionInsight Manager and click **Tenant Resources**. The displayed page integrates multiple functions such as tenant lifecycle management, tenant resource configuration, tenant service association, and tenant resource usage statistics, delivering a mature multi-tenant management model and achieving centralized tenant and service management.
  - Intuitive UI: MRS provides the graphical multi-tenant management interface and manages and operates multiple levels of tenants using the tree structure. Additionally, it integrates the basic information and resource quota of the current tenant in one interface to facilitate O&M and management, as shown in [Figure 6-55](#).

**Figure 6-55** Multi-tenant management (using Manager 3.x as an example)

- Hierarchical tenant management: MRS allows you to add sub-tenants to an existing tenant to re-configure resources. Sub-tenants of level-1 tenants are level-2 tenants. So on and so forth. FusionInsight Manager provides enterprises with a field-tested multi-tenant management model, enabling centralized tenant and service management.
- **Simplified permission management**

MRS hides internal permission management details from common users and simplifies permission management operations for administrators, improving usability and user experience of tenant permission management.

  - MRS employs role-based access control (RBAC) to configure different permissions for users based on service scenarios during multi-tenant management.
  - The administrator of tenants has tenant management permissions, including viewing resources and services of the current tenant, adding or deleting sub-tenants of the current tenant, and managing permissions of sub-tenants' resources. MRS supports setting of the administrator for a single tenant so that the management over this tenant can be delegated to a user who is not the system administrator.
  - Roles of a tenant have all permissions on the computing resources and storage resources of the tenant. When a tenant is created, the system automatically creates roles for this tenant. You can add a user and assign the tenant roles to the user so that the user can use the resources of the tenant.
- **Easy resource management**
  - Free resource configuration

You can configure the compute resources and storage resources during the creation of a tenant and add, modify, or delete the resources of the tenant.

Permissions of the roles that are granted to a tenant are updated automatically when you modify the computing or storage resources of the tenant.

– Resource usage statistics

Resource usage statistics are critical for administrators to determine O&M activities based on the status of cluster applications and services, improving the cluster O&M efficiency. MRS displays the resource statistics of tenants in **Resource Quota**, including the vCores, memory, and HDFS storage resources.

 NOTE

- The available resources of the Capacity scheduler and Superior scheduler are calculated as follows:
  - Capacity
 

Available YARN resources (memory and CPU) = Resource capacity (%) x Total capacity of the resource pool

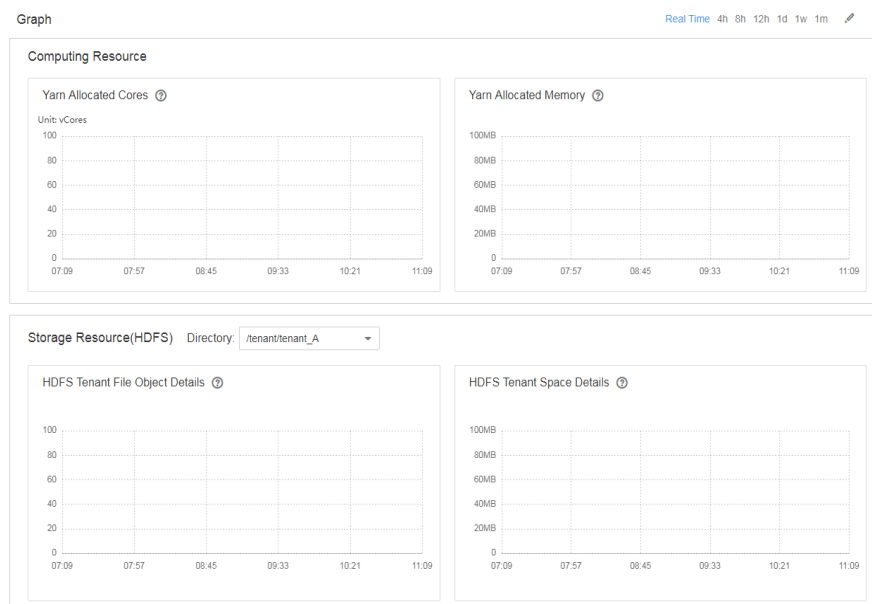
When a queue spans multiple resource pools, its available resources are the combined total of resources allocated by each pool.
  - Superior
 



The available YARN resources (memory and CPU) are allocated in proportion based on the queue weight.
- When the tenant administrator is bound to a tenant role, the tenant administrator has the permissions to manage the tenant and use all resources of the tenant.

– Graphical resource monitoring

Graphical resource monitoring supports the graphical display of monitoring metrics listed in [Table 6-37](#), as shown in [Figure 6-56](#).

**Figure 6-56** Refined monitoring (using Manager 3.x as an example)



By default, real-time monitoring data is displayed. You can click  to customize a time range. Click  and choose **Export** from the shortcut menu to export monitoring information.

**Table 6-37** Monitoring metrics

| Service | Metric Item                                                                                                                                                                     | Description                                                                                                                                                                                                                                                                                |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HDFS    | HDFS Tenant Space Details <ul style="list-style-type: none"> <li>Allocated Space</li> <li>Used Space</li> </ul>                                                                 | HDFS can monitor a specified storage directory. The storage directory is the same as the directory added by the current tenant in <b>Resource</b> .                                                                                                                                        |
|         | HDFS Tenant File Object Details <ul style="list-style-type: none"> <li>Number of Used File Objects</li> </ul>                                                                   |                                                                                                                                                                                                                                                                                            |
| YARN    | Yarn Allocated Cores <ul style="list-style-type: none"> <li>Maximum Number of CPU Cores in an AM</li> <li>Allocated Cores</li> <li>Number of Used CPU Cores in an AM</li> </ul> | Monitoring information of the current tenant is displayed. If no sub-item is configured for a tenant, this information is not displayed.<br><br>The monitoring data is obtained from <b>Scheduler &gt; Application Queues &gt; Queue: <i>Tenant name</i></b> on the native web UI of YARN. |
|         | Yarn Allocated Memory <ul style="list-style-type: none"> <li>Allocated Maximum AM Memory</li> <li>Allocated Memory</li> <li>Used AM Memory</li> </ul>                           |                                                                                                                                                                                                                                                                                            |

## 6.9.2 Using MRS Multi-Tenancy

### Scenarios

Tenants are used in resource control and service isolation scenarios. You need to determine the service scenarios of cluster resources and then plan tenants.

Multi-tenancy involves three types of operations: creating a tenant, managing tenants, and managing resources. [Table 6-38](#) describes these operations.

**Table 6-38** Multi-tenant operations


| Operation          | Action                                                                                                                                                                                                                                                                                                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Creating a tenant  | <ul style="list-style-type: none"> <li>• Creating a tenant</li> <li>• Creating a sub-tenant</li> <li>• Adding a user and assigning a tenant role to the user</li> </ul>                                                                                                                               | <p>During the creation of a tenant, you can configure its computing resources, storage resources, and associated services based on service requirements. In addition, you can add users to the tenant and bind necessary roles to these users.</p> <p>A user to create a level-1 tenant needs to be bound to the <b>Manager_administrator</b> or <b>System_administrator</b> role.</p> <p>A user to create a sub-tenant needs to be bound to the role of the parent tenant at least.</p> |
| Managing tenants   | <ul style="list-style-type: none"> <li>• Managing a tenant directory</li> <li>• Restoring tenant data</li> <li>• Clearing non-associated queues of a tenant</li> <li>• Deleting a tenant</li> </ul>                                                                                                   | <p>You can edit tenants as services change.</p> <p>A user to manage or delete a level-1 tenant or restore tenant data needs to be assigned the <b>Manager_administrator</b> or <b>System_administrator</b> role.</p> <p>A user to manage or delete a sub-tenant needs to be assigned the role of the parent tenant at least.</p>                                                                                                                                                         |
| Managing resources | <ul style="list-style-type: none"> <li>• Creating a resource pool</li> <li>• Modifying a resource pool</li> <li>• Deleting a resource pool</li> <li>• Configuring a queue</li> <li>• Configuring the queue capacity policy of a resource pool</li> <li>• Clearing configuration of a queue</li> </ul> | <p>You can reconfigure resources for tenants as the services change.</p> <p>A user to manage resources needs to be assigned the <b>Manager_administrator</b> or <b>System_administrator</b> role.</p>                                                                                                                                                                                                                                                                                    |

## Creating Tenants and Sub-Tenants

Administrators need to determine the service scenarios of cluster resources and then plan tenants. After that, administrators add tenants and configure dynamic resources, storage resources, and associated services for the tenants on MRS.

**Table 6-39** shows the process for creating a tenant.

**Figure 6-57** Creating a tenant



**Table 6-39** Operations for creating a tenant

| Operation                                               | Description                                                                                                                                                                                                   |
|---------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Creating a tenant                                       | You can configure the compute resources, storage resources, and associated services of the tenant.                                                                                                            |
| Creating a sub-tenant                                   | You can configure the compute resources, storage resources, and associated services of the sub-tenant.                                                                                                        |
| Adding a user and assigning the tenant role to the user | If a user wants to use the resources of tenant <b>tenant1</b> or add or delete sub-tenants for <b>tenant1</b> , the user must be bound to both the <b>Manager_tenant</b> and <b>tenant1_Cluster ID</b> roles. |

## 6.9.3 Configuring MRS Tenants

### 6.9.3.1 Creating an MRS Tenant

You can create a tenant on MRS Manager to specify the resource usage.

#### Prerequisites

- A tenant name has been planned. The name must not be the same as that of a role or Yarn queue that exists in the current cluster.
- If a tenant requires storage resources, a storage directory has been planned based on service requirements, and the planned directory does not exist under the HDFS directory.
- The resources that can be allocated to the current tenant have been planned and the sum of the resource percentages of direct sub-tenants/sub-tenants under the parent tenant at every level does not exceed 100%.
- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

#### Adding a Tenant on the MRS Management Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS cluster details page, click **Tenant**.

**Step 4** Click **Create Tenant**. On the displayed page, configure tenant properties by referring to the following table.

**Table 6-40** Tenant property parameters (MRS 3.x)

| Parameter | Description                                                                                                                                                                                                                                                                                                        |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name      | <ul style="list-style-type: none"><li>• Name of the current tenant, which can contain digits, letters, and underscores (_).</li><li>• Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li></ul> |



| Parameter        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tenant Type      | <p>Whether a tenant is a leaf tenant.</p> <p>In some versions, this parameter is named <b>Tenant Resource Type</b>.</p> <ul style="list-style-type: none"> <li>• When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li> <li>• When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant.</li> </ul> <p><b>NOTE</b><br/>MRS 3.2.0 or later: If you select <b>ClickHouse</b> for <b>Service</b>, this parameter can only be set to <b>Leaf Tenant</b>.</p>                                                                                                                       |
| Compute Resource | <p>Compute resources can be used by the tenant.</p> <ul style="list-style-type: none"> <li>• When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name. <ul style="list-style-type: none"> <li>– A leaf tenant can directly submit jobs to the queue.</li> <li>– A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li> </ul> </li> <li>• If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li> </ul> |

| Parameter                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuration Mode                  | <p>If <b>Yarn</b> is selected for <b>Compute Resource</b>, this parameter can be set to <b>Basic</b> or <b>Advanced</b>.</p> <ul style="list-style-type: none"><li>• <b>Basic</b>: Configure the percentage of compute resources used by the tenant in the default resource pool by specifying <b>Default Resource Pool Capacity (%)</b>.</li><li>• <b>Advanced</b>: Configure the following parameters for advanced settings:<ul style="list-style-type: none"><li>– <b>Weight</b>: Tenant resource weight. The value ranges from 0 to 100. Tenant resource weight = Tenant weight/Total weight of tenants at the same level</li><li>– <b>Minimum Resources</b>: resources preempted by the tenant. The value is a percentage or absolute value of the parent tenant's resources. When a tenant's workload is light, their resources are automatically lent to other tenants. When available resources are fewer than <b>Minimum Resources</b>, the tenant can preempt the resources that were lent out.</li><li>– <b>Maximum Resources</b>: maximum resources that can be used by a tenant. The value is a percentage or absolute value of the parent tenant's resources.</li><li>– <b>Reserved Resources</b>: resources reserved for the tenant. The reserved resources cannot be used by other tenants even if no job is using resources of the tenant. The value is a percentage or absolute value of the parent tenant's resources.</li></ul></li></ul> |
| Storage Resource                    | <p>Storage resources selected for the current tenant.</p> <ul style="list-style-type: none"><li>• If <b>HDFS</b> is selected, the system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory. When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory.</li><li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Maximum Number of Files/Directories | <p>Maximum number of files or directories that can be created in HDFS. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Space Quota | <p>Specifies the quota for HDFS storage space used by the current tenant. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p> <ul style="list-style-type: none"> <li>• The minimum value is <b>1</b>, and the unit is MB or GB.</li> <li>• This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li> <li>• If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li> </ul> <p><b>NOTE</b><br/>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500</b> MB, the actual space for storing files is about 250 MB (500/2 = 250).</p> |
| Storage Path        | <p>HDFS directory for tenant resource data.</p> <ul style="list-style-type: none"> <li>• The system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory by default. For example, the default HDFS storage directory for tenant <b>ta1</b> is <b>/tenant/ta1</b>.</li> <li>• When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory. The storage path is customizable.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Service             | <p>Other service resources associated with the tenant. (This parameter is grayed out if the cluster has no components that support service association.)</p> <p>Click <b>Associate Service</b> and select a service name from the <b>Service</b> drop-down list box. If <b>Association Mode</b> is set to <b>Exclusive</b>, service resources are occupied exclusively. If <b>share</b> is selected, service resources are shared.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description         | Specifies the description of the current tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Table 6-41** Tenant parameters (MRS 2.x and earlier versions)

| Parameter                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                                    | <ul style="list-style-type: none"><li>Name of the current tenant, which can contain digits, letters, and underscores (_).</li><li>Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Tenant Type                             | Whether the tenant is a leaf tenant. <ul style="list-style-type: none"><li>When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li><li>When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Compute Resource                        | Compute resources can be used by the tenant. (In some versions, the parameter is named <b>Dynamic Resource</b> .) <ul style="list-style-type: none"><li>When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name.<ul style="list-style-type: none"><li>A leaf tenant can directly submit jobs to the queue.</li><li>A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li></ul></li><li>If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li></ul> |
| Default Resource Pool Capacity (%)      | Percentage of the compute resources used by the current tenant in the <b>default</b> resource pool. This parameter is configured when YARN is the <b>Compute Resource</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Default Resource Pool Max. Capacity (%) | Maximum percentage of the compute resources used by the current tenant in the <b>default</b> resource pool. This parameter is configured when YARN is the <b>Compute Resource</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Parameter                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Resource         | <p>Storage resources selected for the current tenant.</p> <ul style="list-style-type: none"><li>• If <b>HDFS</b> is selected, the system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory. When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory.</li><li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Storage Space Quota (MB) | <p>Specifies the quota for HDFS storage space used by the current tenant. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p> <ul style="list-style-type: none"><li>• The minimum value is <b>1</b>, and the unit is MB or GB.</li><li>• This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li><li>• If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li></ul> <p><b>NOTE</b></p> <p>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500</b> MB, the actual space for storing files is about 250 MB (<math>500/2 = 250</math>).</p> |
| Storage Path             | <p>HDFS storage directory for the tenant.</p> <ul style="list-style-type: none"><li>• The system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory by default. For example, the default HDFS storage directory for tenant <b>ta1</b> is <b>/tenant/ta1</b>.</li><li>• When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory. The storage path is customizable.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Service                  | <p>Other service resources associated with the tenant. (This parameter is grayed out if the cluster has no components that support service association.)</p> <p>Click <b>Associate Service</b> and select a service name from the <b>Service</b> drop-down list box. If <b>Association Mode</b> is set to <b>Exclusive</b>, service resources are occupied exclusively. If <b>share</b> is selected, service resources are shared.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Description              | Specifies the description of the current tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Step 5** Click **OK** to save. Wait until the tenant is created.

In the tenant list on the left, click the name of the created tenant to view its basic information, resource quotas, and charts.


 **NOTE**

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. The role and its permissions are controlled by the system automatically and cannot be controlled manually under **Manage Role**.
- If you want to use the tenant, create a system user and assign the Manager\_tenant role and the role corresponding to the tenant to the user. For details, see [Binding Tenant to an MRS Cluster User](#).

----End

## Adding a Tenant on Manager (for MRS 3.x and Later)

**Step 1** Log in to Manager and choose **Tenant Resources**.

**Step 2** Click . On the page that is displayed, configure tenant properties according to [Table 6-42](#).

**Table 6-42** Tenant property parameters (MRS 3.x)

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name        | <ul style="list-style-type: none"><li>• Name of the current tenant, which can contain digits, letters, and underscores (_).</li><li>• Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li></ul>                                                                                                                                                                                                                                                                     |
| Tenant Type | <p>Whether a tenant is a leaf tenant.</p> <p>In some versions, this parameter is named <b>Tenant Resource Type</b>.</p> <ul style="list-style-type: none"><li>• When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li><li>• When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant.</li></ul> <p><b>NOTE</b><br/>MRS 3.2.0 or later: If you select <b>ClickHouse</b> for <b>Service</b>, this parameter can only be set to <b>Leaf Tenant</b>.</p> |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Compute Resource   | <p>Compute resources can be used by the tenant.</p> <ul style="list-style-type: none"> <li>• When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name. <ul style="list-style-type: none"> <li>– A leaf tenant can directly submit jobs to the queue.</li> <li>– A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li> </ul> </li> <li>• If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Configuration Mode | <p>If <b>Yarn</b> is selected for <b>Compute Resource</b>, this parameter can be set to <b>Basic</b> or <b>Advanced</b>.</p> <ul style="list-style-type: none"> <li>• <b>Basic</b>: Configure the percentage of compute resources used by the tenant in the default resource pool by specifying <b>Default Resource Pool Capacity (%)</b>.</li> <li>• <b>Advanced</b>: Configure the following parameters for advanced settings: <ul style="list-style-type: none"> <li>– <b>Weight</b>: Tenant resource weight. The value ranges from 0 to 100. Tenant resource weight = Tenant weight/Total weight of tenants at the same level</li> <li>– <b>Minimum Resources</b>: resources preempted by the tenant. The value is a percentage or absolute value of the parent tenant's resources. When a tenant's workload is light, their resources are automatically lent to other tenants. When available resources are fewer than <b>Minimum Resources</b>, the tenant can preempt the resources that were lent out.</li> <li>– <b>Maximum Resources</b>: maximum resources that can be used by a tenant. The value is a percentage or absolute value of the parent tenant's resources.</li> <li>– <b>Reserved Resources</b>: resources reserved for the tenant. The reserved resources cannot be used by other tenants even if no job is using resources of the tenant. The value is a percentage or absolute value of the parent tenant's resources.</li> </ul> </li> </ul> |

| Parameter                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Resource                    | <p>Storage resources selected for the current tenant.</p> <ul style="list-style-type: none"><li>• If <b>HDFS</b> is selected, the system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory. When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory.</li><li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Maximum Number of Files/Directories | <p>Maximum number of files or directories that can be created in HDFS. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Storage Space Quota                 | <p>Specifies the quota for HDFS storage space used by the current tenant. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p> <ul style="list-style-type: none"><li>• The minimum value is <b>1</b>, and the unit is MB or GB.</li><li>• This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li><li>• If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li></ul> <p><b>NOTE</b><br/>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500 MB</b>, the actual space for storing files is about 250 MB (<math>500/2 = 250</math>).</p> |
| Storage Path                        | <p>HDFS directory for tenant resource data.</p> <ul style="list-style-type: none"><li>• The system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory by default. For example, the default HDFS storage directory for tenant <b>ta1</b> is <b>/tenant/ta1</b>.</li><li>• When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory. The storage path is customizable.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                          |



| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Service     | Other service resources associated with the tenant. (This parameter is grayed out if the cluster has no components that support service association.)<br>Click <b>Associate Service</b> and select a service name from the <b>Service</b> drop-down list box. If <b>Association Mode</b> is set to <b>Exclusive</b> , service resources are occupied exclusively. If <b>share</b> is selected, service resources are shared. |
| Description | Specifies the description of the current tenant.                                                                                                                                                                                                                                                                                                                                                                             |

 NOTE

Roles, computing resources, and storage resources are automatically created when tenants are created.

- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System > Permission > Role**. The role name is in the format of *Tenant name\_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and bind the user to the role of the tenant. For details, see [Binding Tenant to an MRS Cluster User](#).
- During the tenant creation, the system automatically creates a YARN queue named after the tenant. If the queue name already exists, the new queue is named **Tenant name-N**. *N* indicates a natural number starting from 1. When a same name exists, the value *N* increases automatically to differentiate the queue from others. For example, **saletenant**, **saletenant-1**, and **saletenant-2**.

If you need to associate a service, click **Associate Service**, configure other service resources associated with the tenant by referring to the following instructions, and click **OK**.

- Set **Service** to **HBase** and **Association Type** to **Exclusive** or **Shared**.

 NOTE

- **Exclusive** indicates that the service resources are used by the tenant exclusively and cannot be associated with other tenants.
- **Shared** indicates that the service resources can be shared with other tenants.
- MRS 3.2.0 or later: Select **ClickHouse** for **Service**.
  - **Association Type**: When **Service** is set to **ClickHouse**, **Association Type** can only be set to **Shared**.
  - **Associate Logical Cluster**: If the logical cluster function is not enabled for ClickHouse, **default\_cluster** is selected by default. If the function is enabled, select the logical cluster to which you want to associate.
  - **CPU Priority**: The CPU priority ranges from -20 to 19. This value is associated with the NICE value of the OS. A smaller value indicates a higher CPU priority.
  - **Memory**: The maximum value of this parameter is **100**, in percentage. For example, if this parameter is set to **80**, the total memory that can be used by the current tenant is calculated as follows: Available memory x 80%.

 NOTE

- Only HBase can be associated with a new tenant. However, HDFS, HBase, and YARN can be associated with existing tenants.
- To associate an existing tenant with service resources, click the target tenant in the tenant list, switch to the **Service Associations** page, and click **Associate Service** to configure resources to be associated with the tenant.
- To disassociate an existing tenant from service resources, click the target tenant in the tenant list, switch to the **Service Associations** page, and click **Delete** in the **Operation** column. In the displayed dialog box, select **I have read the information and understand the impact** and click **OK**.

**Step 3** Click **OK**. Wait until the system displays a message indicating that the tenant is successfully created.

----End

## Adding a Tenant on Manager (for MRS 2.x and Later)

**Step 1** On MRS Manager, click **Tenant**.

**Step 2** Click **Create Tenant**. On the displayed page, configure tenant properties.

**Table 6-43** Tenant parameters (MRS 2.x and earlier versions)

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                    |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name        | <ul style="list-style-type: none"><li>• Name of the current tenant, which can contain digits, letters, and underscores (_).</li><li>• Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li></ul>                             |
| Tenant Type | Whether the tenant is a leaf tenant. <ul style="list-style-type: none"><li>• When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li><li>• When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant.</li></ul> |

| Parameter                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Compute Resource                        | <p>Compute resources can be used by the tenant. (In some versions, the parameter is named <b>Dynamic Resource</b>.)</p> <ul style="list-style-type: none"><li>• When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name.<ul style="list-style-type: none"><li>– A leaf tenant can directly submit jobs to the queue.</li><li>– A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li></ul></li><li>• If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li></ul> |
| Default Resource Pool Capacity (%)      | <p>Percentage of the compute resources used by the current tenant in the <b>default</b> resource pool. This parameter is configured when YARN is the <b>Compute Resource</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Default Resource Pool Max. Capacity (%) | <p>Maximum percentage of the compute resources used by the current tenant in the <b>default</b> resource pool. This parameter is configured when YARN is the <b>Compute Resource</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Storage Resource                        | <p>Storage resources selected for the current tenant.</p> <ul style="list-style-type: none"><li>• If <b>HDFS</b> is selected, the system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory. When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory.</li><li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS.</li></ul>                                                                                                                                                                                                                                                       |

| Parameter                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Space Quota (MB) | <p>Specifies the quota for HDFS storage space used by the current tenant. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p> <ul style="list-style-type: none"><li>• The minimum value is <b>1</b>, and the unit is MB or GB.</li><li>• This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li><li>• If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li></ul> <p><b>NOTE</b><br/>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500</b> MB, the actual space for storing files is about 250 MB (<math>500/2 = 250</math>).</p> |
| Storage Path             | <p>HDFS storage directory for the tenant.</p> <ul style="list-style-type: none"><li>• The system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory by default. For example, the default HDFS storage directory for tenant <b>ta1</b> is <b>/tenant/ta1</b>.</li><li>• When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory. The storage path is customizable.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Service                  | <p>Other service resources associated with the tenant. (This parameter is grayed out if the cluster has no components that support service association.)</p> <p>Click <b>Associate Service</b> and select a service name from the <b>Service</b> drop-down list box. If <b>Association Mode</b> is set to <b>Exclusive</b>, service resources are occupied exclusively. If <b>share</b> is selected, service resources are shared.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Description              | Specifies the description of the current tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Step 3** Click **OK** to save. Wait until the tenant is created.

 NOTE

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. The role and its permissions are controlled by the system automatically and cannot be controlled manually under **Manage Role**.
- If you want to use the tenant, create a system user and assign the `Manager_tenant` role and the role corresponding to the tenant to the user. For details, see [Binding Tenant to an MRS Cluster User](#).

----End

### 6.9.3.2 Creating an MRS Sub-Tenant

You can create sub-tenants on MRS and allocate resources of the current tenant to the sub-tenants based on the resource consumption and isolation planning and requirements of services.

If the tenant is a non-leaf tenant, you can add sub-tenants. Sub-tenants cannot be created for a leaf tenant.

#### Prerequisites

- A non-leaf tenant has been added by referring to [Creating an MRS Tenant](#).
- A tenant name has been planned. The name must not be the same as that of a role or Yarn queue that exists in the current cluster.
- If a sub-tenant requires storage resources, a storage directory has been planned based on service requirements, and the planned directory does not exist under the storage directory of the parent tenant.
- The resources that can be allocated to the current tenant have been planned and the sum of the resource percentages of direct sub-tenants under the parent tenant at every level does not exceed 100%.
- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Adding an MRS Sub-Tenant on the MRS Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS details page, click **Tenant**.

**Step 4** In the tenant list on the left, move the cursor to the tenant node to which a sub-tenant is to be added. Click **Create sub-tenant**. On the displayed page, configure the sub-tenant attributes according to the following table:

**Table 6-44** Tenant property parameters (MRS 3.x)

| Parameter        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name             | <ul style="list-style-type: none"><li>• Name of the current tenant, which can contain digits, letters, and underscores (_).</li><li>• Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                      |
| Tenant Type      | <p>Whether a tenant is a leaf tenant.</p> <p>In some versions, this parameter is named <b>Tenant Resource Type</b>.</p> <ul style="list-style-type: none"><li>• When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li><li>• When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant.</li></ul> <p><b>NOTE</b><br/>MRS 3.2.0 or later: If you select <b>ClickHouse</b> for <b>Service</b>, this parameter can only be set to <b>Leaf Tenant</b>.</p>                                                                                                                  |
| Compute Resource | <p>Compute resources can be used by the tenant.</p> <ul style="list-style-type: none"><li>• When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name.<ul style="list-style-type: none"><li>– A leaf tenant can directly submit jobs to the queue.</li><li>– A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li></ul></li><li>• If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li></ul> |

| Parameter                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuration Mode                  | <p>If <b>Yarn</b> is selected for <b>Compute Resource</b>, this parameter can be set to <b>Basic</b> or <b>Advanced</b>.</p> <ul style="list-style-type: none"> <li>• <b>Basic:</b> Configure the percentage of compute resources used by the tenant in the default resource pool by specifying <b>Default Resource Pool Capacity (%)</b>.</li> <li>• <b>Advanced:</b> Configure the following parameters for advanced settings: <ul style="list-style-type: none"> <li>– <b>Weight:</b> Tenant resource weight. The value ranges from 0 to 100. Tenant resource weight = Tenant weight/Total weight of tenants at the same level</li> <li>– <b>Minimum Resources:</b> resources preempted by the tenant. The value is a percentage or absolute value of the parent tenant's resources. When a tenant's workload is light, their resources are automatically lent to other tenants. When available resources are fewer than <b>Minimum Resources</b>, the tenant can preempt the resources that were lent out.</li> <li>– <b>Maximum Resources:</b> maximum resources that can be used by a tenant. The value is a percentage or absolute value of the parent tenant's resources.</li> <li>– <b>Reserved Resources:</b> resources reserved for the tenant. The reserved resources cannot be used by other tenants even if no job is using resources of the tenant. The value is a percentage or absolute value of the parent tenant's resources.</li> </ul> </li> </ul> |
| Storage Resource                    | <p>Storage resources selected for the current tenant.</p> <ul style="list-style-type: none"> <li>• If <b>HDFS</b> is selected, the system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory. When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory.</li> <li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Maximum Number of Files/Directories | <p>Maximum number of files or directories that can be created in HDFS. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Space Quota | <p>Specifies the quota for HDFS storage space used by the current tenant. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p> <ul style="list-style-type: none"> <li>• The minimum value is <b>1</b>, and the unit is MB or GB.</li> <li>• This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li> <li>• If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li> </ul> <p><b>NOTE</b><br/>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500</b> MB, the actual space for storing files is about 250 MB (500/2 = 250).</p> |
| Storage Path        | <p>HDFS directory for tenant resource data.</p> <ul style="list-style-type: none"> <li>• The system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory by default. For example, the default HDFS storage directory for tenant <b>ta1</b> is <b>/tenant/ta1</b>.</li> <li>• When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory. The storage path is customizable.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Service             | <p>Other service resources associated with the tenant. (This parameter is grayed out if the cluster has no components that support service association.)</p> <p>Click <b>Associate Service</b> and select a service name from the <b>Service</b> drop-down list box. If <b>Association Mode</b> is set to <b>Exclusive</b>, service resources are occupied exclusively. If <b>share</b> is selected, service resources are shared.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description         | Specifies the description of the current tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



**Table 6-45** Tenant parameters (MRS 2.x and earlier versions)

| Parameter                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                                    | <ul style="list-style-type: none"><li>Name of the current tenant, which can contain digits, letters, and underscores (_).</li><li>Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Tenant Type                             | Whether the tenant is a leaf tenant. <ul style="list-style-type: none"><li>When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li><li>When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Compute Resource                        | Compute resources can be used by the tenant. (In some versions, the parameter is named <b>Dynamic Resource</b> .) <ul style="list-style-type: none"><li>When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name.<ul style="list-style-type: none"><li>A leaf tenant can directly submit jobs to the queue.</li><li>A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li></ul></li><li>If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li></ul> |
| Default Resource Pool Capacity (%)      | Percentage of the compute resources used by the current tenant in the <b>default</b> resource pool. This parameter is configured when YARN is the <b>Compute Resource</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Default Resource Pool Max. Capacity (%) | Maximum percentage of the compute resources used by the current tenant in the <b>default</b> resource pool. This parameter is configured when YARN is the <b>Compute Resource</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Parameter                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Resource         | <p>Storage resources selected for the current tenant.</p> <ul style="list-style-type: none"><li>• If <b>HDFS</b> is selected, the system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory. When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory.</li><li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Storage Space Quota (MB) | <p>Specifies the quota for HDFS storage space used by the current tenant. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p> <ul style="list-style-type: none"><li>• The minimum value is <b>1</b>, and the unit is MB or GB.</li><li>• This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li><li>• If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li></ul> <p><b>NOTE</b></p> <p>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500</b> MB, the actual space for storing files is about 250 MB (<math>500/2 = 250</math>).</p> |
| Storage Path             | <p>HDFS storage directory for the tenant.</p> <ul style="list-style-type: none"><li>• The system automatically creates a file folder named after the tenant name in the <b>/tenant</b> directory by default. For example, the default HDFS storage directory for tenant <b>ta1</b> is <b>/tenant/ta1</b>.</li><li>• When a tenant is created for the first time, the system automatically creates the <b>/tenant</b> directory in the HDFS root directory. The storage path is customizable.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Service                  | <p>Other service resources associated with the tenant. (This parameter is grayed out if the cluster has no components that support service association.)</p> <p>Click <b>Associate Service</b> and select a service name from the <b>Service</b> drop-down list box. If <b>Association Mode</b> is set to <b>Exclusive</b>, service resources are occupied exclusively. If <b>share</b> is selected, service resources are shared.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Description              | Specifies the description of the current tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Step 5** Click **OK** to save the settings.

It takes a few minutes to save the settings. If the **Tenant created successfully** is displayed in the upper-right corner, the tenant is added successfully. The tenant is created successfully.


 **NOTE**

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System > Permission > Role**. The role name is in the format of *Tenant name\_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and assign the tenant role to the user. For details, see [Binding Tenant to an MRS Cluster User](#).
- The sub-tenant can further allocate the resources of its parent tenant. The sum of the resource percentages of direct sub-tenants under a parent tenant at each level cannot exceed 100%. The sum of the computing resource percentages of all level-1 tenants cannot exceed 100%.

----End

## Adding a Sub-Tenant on Manager (for MRS 3.x and Later)

**Step 1** Log in to FusionInsight Manager and choose **Tenant Resources**.

**Step 2** In the tenant list on the left, select a parent tenant and click . On the displayed page, set the parameters.

- To use a Superior scheduler cluster, configure sub-tenant properties by referring to [Table 6-46](#).
- To use the Capacity scheduler cluster, configure sub-tenant properties by referring to [Table 6-47](#).

 **NOTE**

- YARN in a new cluster uses the Superior scheduler by default. You can also switch the scheduler by referring to [Switching the MRS Tenant Resource Scheduler](#).
- To query the scheduler type, log in to Manager and search for the **yarn.resourcemanager.scheduler.class** parameter on the **All Configurations** page of the YARN service.

**Table 6-46** Sub-tenant parameters (Superior scheduler)

| Parameter              | Description                                                                                  |
|------------------------|----------------------------------------------------------------------------------------------|
| Cluster                | Cluster to which the parent tenant belongs.                                                  |
| Parent Tenant Resource | Name of the parent tenant. (In some versions, the parameter is named <b>Parent Tenant</b> .) |

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                 | <ul style="list-style-type: none"> <li>• Name of the current tenant, which can contain digits, letters, and underscores (_).</li> <li>• Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                |
| Tenant Resource Type | <p>Whether a tenant is a leaf tenant. (In some versions, the parameter is named <b>Tenant Type</b>.)</p> <ul style="list-style-type: none"> <li>• When <b>Leaf Tenant Resource</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added. (In some versions, the parameter is named <b>Leaf Tenant</b>.)</li> <li>• When <b>Non-leaf Tenant Resource</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant. However, the tenant depth cannot exceed 5 levels. (In some versions, the parameter is named <b>Non-Leaf Tenant</b>.)</li> </ul>                                                                                          |
| Compute Resource     | <p>Dynamic compute resources for the current tenant.</p> <ul style="list-style-type: none"> <li>• When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name. <ul style="list-style-type: none"> <li>– A leaf tenant can directly submit jobs to the queue.</li> <li>– A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li> </ul> </li> <li>• If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li> </ul> |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuration Mode | <p>Configuration mode of compute resource parameters. If <b>Yarn</b> is selected for <b>Compute Resource</b>, this parameter can be set to <b>Basic</b> or <b>Advanced</b>.</p> <ul style="list-style-type: none"> <li>• <b>Basic</b>: Configure the percentage of compute resources used by the tenant in the default resource pool by specifying <b>Default Resource Pool Capacity (%)</b>.</li> <li>• <b>Advanced</b>: Configure the following parameters for advanced settings: <ul style="list-style-type: none"> <li>- <b>Weight</b>: Tenant resource weight. The value ranges from 0 to 100. Tenant resource weight = Tenant weight/Total weight of tenants at the same level</li> <li>- <b>Minimum Resources</b>: resources preempted by the tenant. The value is a percentage or absolute value of the parent tenant's resources. When a tenant's workload is light, their resources are automatically lent to other tenants. When available resources are fewer than <b>Minimum Resources</b>, the tenant can preempt the resources that were lent out.</li> <li>- <b>Maximum Resources</b>: maximum resources that can be used by a tenant. The value is a percentage or absolute value of the parent tenant's resources.</li> <li>- <b>Reserved Resources</b>: resources reserved for the tenant. The reserved resources cannot be used by other tenants even if no job is using resources of the tenant. The value is a percentage or absolute value of the parent tenant's resources.</li> </ul> </li> </ul> |
| Storage Resource   | <p>Storage resources selected for the current tenant.</p> <ul style="list-style-type: none"> <li>• When <b>HDFS</b> is selected, the system automatically creates a folder named after the sub-tenant in the HDFS parent tenant directory.</li> <li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS. If the parent tenant does not have storage resources, the sub-tenant cannot use storage resources.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Quota              | <p>Quota for files and directories. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Space Quota | <p>Quota for HDFS storage space used by the current tenant. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p> <ul style="list-style-type: none"> <li>The minimum value is <b>1</b>, and the maximum value is the total storage quota of the parent tenant. The unit is MB or GB.</li> <li>This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li> <li>If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li> <li>If this quota is greater than the quota of the parent tenant, the actual storage space does not exceed the quota of the parent tenant.</li> </ul> <p><b>NOTE</b><br/>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500</b> MB, the actual space for storing files is about 250 MB (<math>500/2 = 250</math>).</p> |
| Storage Path        | <p>Specifies the tenant's HDFS storage directory.</p> <ul style="list-style-type: none"> <li>The system automatically creates a file folder named after the sub-tenant name in the directory of the parent tenant by default. For example, if the sub-tenant is <b>ta1s</b> and the parent directory is <b>tenant/ta1</b>, the system sets this parameter for the sub-tenant to <b>tenant/ta1/ta1s</b>.</li> <li>The storage path is customizable in the parent directory. The parent directory for the storage path must be the storage directory of the parent tenant.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Service             | Whether to associate resources of other services. For details, see <a href="#">Step 4</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Description         | Description of the tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 6-47** Sub-tenant parameters (Capacity scheduler)

| Parameter              | Description                                                                                  |
|------------------------|----------------------------------------------------------------------------------------------|
| Cluster                | Cluster to which the parent tenant belongs.                                                  |
| Parent Tenant Resource | Name of the parent tenant. (In some versions, the parameter is named <b>Parent Tenant</b> .) |

| Parameter                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                                   | <ul style="list-style-type: none"><li>• Name of the current tenant, which can contain digits, letters, and underscores (_).</li><li>• Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                    |
| Tenant Type                            | Whether a tenant is a leaf tenant. <ul style="list-style-type: none"><li>• When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li><li>• When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant.</li></ul>                                                                                                                                                                                                                                                                                                                                                          |
| Compute Resource                       | Dynamic compute resources for the current tenant. <ul style="list-style-type: none"><li>• When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name.<ul style="list-style-type: none"><li>- A leaf tenant can directly submit jobs to the queue.</li><li>- A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li></ul></li><li>• If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li></ul> |
| Default Resource Pool Capacity (%)     | Percentage of computing resources used by the current tenant. The base value is the total resources of the parent tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Default Resource Pool Max Capacity (%) | Specifies the maximum percentage of the computing resources used by the current tenant. The base value is the total resources of the parent tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Storage Resource                       | Storage resources selected for the current tenant. <ul style="list-style-type: none"><li>• When <b>HDFS</b> is selected, the system automatically creates a folder named after the sub-tenant in the HDFS parent tenant directory.</li><li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS. If the parent tenant does not have storage resources, the sub-tenant cannot use storage resources.</li></ul>                                                                                                                                                                                                                            |
| Quota                                  | Quota for files and directories. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Space Quota | <p>Quota for HDFS storage space used by the current tenant. Set this parameter when <b>Storage Resource</b> is <b>HDFS</b>.</p> <ul style="list-style-type: none"><li>• The minimum value is <b>1</b>, and the maximum value is the total storage quota of the parent tenant. The unit is MB or GB.</li><li>• This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li><li>• If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li><li>• If this quota is greater than the quota of the parent tenant, the actual storage space does not exceed the quota of the parent tenant.</li></ul> <p><b>NOTE</b><br/>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500 MB</b>, the actual space for storing files is about 250 MB (500/2 = 250).</p> |
| Storage Path        | <p>Specifies the tenant's HDFS storage directory.</p> <ul style="list-style-type: none"><li>• The system automatically creates a file folder named after the sub-tenant name in the directory of the parent tenant by default. For example, if the sub-tenant is <b>ta1s</b> and the parent directory is <b>tenant/ta1</b>, the system sets this parameter for the sub-tenant to <b>tenant/ta1/ta1s</b>.</li><li>• The storage path is customizable in the parent directory. The parent directory for the storage path must be the storage directory of the parent tenant.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Description         | Description of the current tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



 NOTE

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System** > **Permission** > **Role**. The role name is in the format of *Tenant name\_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and assign the tenant role to the user. For details, see [Binding Tenant to an MRS Cluster User](#).
- The sub-tenant can further allocate the resources of its parent tenant. The sum of the resource percentages of direct sub-tenants under a parent tenant at each level cannot exceed 100%. The sum of the computing resource percentages of all level-1 tenants cannot exceed 100%.

**Step 3** Check whether the current tenant needs to be associated with resources of other services.

- If yes, go to [Step 4](#).
- If no, go to [Step 5](#).

**Step 4** Click **Associate Service** to configure other service resources used by the current tenant.

1. Set **Services** to **HBase**.
2. Set **Association Type** as follows:
  - **Exclusive** indicates that the service resources are used by the tenant exclusively and cannot be associated with other tenants.
  - **Shared** indicates that the service resources can be shared with other tenants.

 NOTE

- Only HBase can be associated with a new tenant. However, HDFS, HBase, and YARN can be associated with existing tenants.
- To associate an existing tenant with service resources, click the target tenant in the tenant list, switch to the **Service Associations** page, and click **Associate Service** to configure resources to be associated with the tenant.
- To disassociate an existing tenant from service resources, click the target tenant in the tenant list, switch to the **Service Associations** page, and click **Delete** in the **Operation** column. In the displayed dialog box, select **I have read the information and understand the impact** and click **OK**.

3. Click **OK**.

**Step 5** Click **OK**. Wait until the system displays a message indicating that the tenant is successfully created.

----End

## Adding a Sub-Tenant on Manager (for MRS 2.x and Later)

**Step 1** On MRS Manager, click **Tenant**.

**Step 2** In the tenant list on the left, move the cursor to the tenant node to which a sub-tenant is to be added. Click **Create sub-tenant**. On the displayed page, configure the sub-tenant attributes according to the following table:

**Table 6-48** Sub-tenant parameters

| Parameter                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Parent tenant                           | Specifies the name of the parent tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Name                                    | <ul style="list-style-type: none"><li>• Name of the current tenant, which can contain digits, letters, and underscores (_).</li><li>• Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or YARN queue that exists in the current cluster.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                    |
| Tenant Type                             | Whether the tenant is a leaf tenant. <ul style="list-style-type: none"><li>• When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li><li>• When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant.</li></ul>                                                                                                                                                                                                                                                                                                                                                        |
| Dynamic Resource                        | Dynamic compute resources for the current tenant. <ul style="list-style-type: none"><li>• When <b>Yarn</b> is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name.<ul style="list-style-type: none"><li>– A leaf tenant can directly submit jobs to the queue.</li><li>– A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li></ul></li><li>• If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li></ul> |
| Default Resource Pool Capacity (%)      | Specifies the percentage of the resources used by the current tenant. The base value is the total resources of the parent tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Default Resource Pool Max. Capacity (%) | Specifies the maximum percentage of the computing resources used by the current tenant. The base value is the total resources of the parent tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Storage Resource                        | Storage resources selected for the current tenant. <ul style="list-style-type: none"><li>• When <b>HDFS</b> is selected, the system automatically creates a folder named after the sub-tenant in the HDFS parent tenant directory.</li><li>• If <b>HDFS</b> is not selected, the system does not create a storage directory under the root directory of HDFS. If the parent tenant does not have storage resources, the sub-tenant cannot use storage resources.</li></ul>                                                                                                                                                                                                                            |

| Parameter                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Space Quota (MB) | <p>Quota for HDFS storage space used by the current tenant.</p> <ul style="list-style-type: none"><li>• The minimum value is <b>1</b>, and the maximum value is the total storage quota of the parent tenant. The unit is MB or GB.</li><li>• This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used.</li><li>• If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li><li>• If this quota is greater than the quota of the parent tenant, the actual storage space does not exceed the quota of the parent tenant.</li></ul> <p><b>NOTE</b><br/>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to <b>500</b> MB, the actual space for storing files is about 250 MB (<math>500/2 = 250</math>).</p> |
| Storage Path             | <p>Specifies the tenant's HDFS storage directory.</p> <ul style="list-style-type: none"><li>• The system automatically creates a file folder named after the sub-tenant name in the directory of the parent tenant by default. For example, if the sub-tenant is <b>ta1s</b> and the parent directory is <b>tenant/ta1</b>, the system sets this parameter for the sub-tenant to <b>tenant/ta1/ta1s</b>.</li><li>• The storage path is customizable in the parent directory. The parent directory for the storage path must be the storage directory of the parent tenant.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Service                  | <p>Click <b>Associate Service</b> and select a service name from the <b>Service</b> drop-down list box. If <b>Association Mode</b> is set to <b>Exclusive</b>, service resources are occupied exclusively. If <b>share</b> is selected, service resources are shared.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Description              | <p>Specifies the description of the current tenant.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**Step 3** Click **OK** to save the settings.

It takes a few minutes to save the settings. If the **Tenant created successfully** is displayed in the upper-right corner, the tenant is added successfully. The tenant is created successfully.

 NOTE

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System** > **Permission** > **Role**. The role name is in the format of *Tenant name\_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and assign the tenant role to the user. For details, see [Binding Tenant to an MRS Cluster User](#).
- The sub-tenant can further allocate the resources of its parent tenant. The sum of the resource percentages of direct sub-tenants under a parent tenant at each level cannot exceed 100%. The sum of the computing resource percentages of all level-1 tenants cannot exceed 100%.

----End

### 6.9.3.3 Binding Tenant to an MRS Cluster User

#### Scenario

A newly created tenant cannot directly log in to the cluster to access resources. You need to add a user for the tenant on FusionInsight Manager and bind the user to the role of the tenant to assign operation permissions to the user.

#### Prerequisites

You have clarified service requirements and created a tenant.

#### Procedure

**Step 1** Log in to FusionInsight Manager and choose **System** > **Permission** > **User**.

**Step 2** If you want to add a user to the system, click **Create**.

**Figure 6-58** Adding a user

User > Create

---

\* Username:

\* User Type:  Human-Machine  
 Machine-Machine

\* Password:

\* Confirm Password:

User Group: [Add](#) [Clear All](#) [Create User Group](#)

Primary Group:

Role: [Add](#) [Clear All](#) [Create Role](#)

Description:

If you want to bind tenant roles to an existing user in the system, locate the row of the user and click **Modify** in the **Operation** column.

Set user attributes according to [Table 6-49](#).

**Table 6-49** User parameters

| Parameter        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Username         | <p>Indicates the current username. The value contains 3 to 32 characters, including digits, letters, underscores (_), hyphens (-), and spaces.</p> <ul style="list-style-type: none"> <li>The username cannot be the same as the OS username of any node in the cluster. Otherwise, the user cannot be used.</li> <li>A username that differs only in alphabetic case from an existing username is not allowed. For example, if <b>User1</b> has been created, you cannot create <b>user1</b>. Enter the correct username when using <b>User1</b>.</li> </ul> |
| User Type        | <p>The options are <b>Human-Machine</b> and <b>Machine-Machine</b>.</p> <ul style="list-style-type: none"> <li><b>Human-Machine</b> user: used for FusionInsight Manager O&amp;M and component client operations. If you select this option, set both <b>Password</b> and <b>Confirm Password</b> accordingly.</li> <li><b>Machine-Machine</b> user: used for application development. If you select this option, the password is randomly generated.</li> </ul>                                                                                              |
| Password         | <p>This parameter is mandatory if <b>User Type</b> is set to <b>Human-Machine</b>.</p> <p>The password must contain 8 to 64 characters of at least four types of the following: uppercase letters, lowercase letters, digits, special characters, and spaces. The password cannot be the username or the username spelled backwards.</p>                                                                                                                                                                                                                      |
| Confirm Password | Enter the password again.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| User Group       | <p>In the <b>User Group</b> area, click <b>Add</b> and select user groups to add the user to the groups.</p> <ul style="list-style-type: none"> <li>If roles have been added to the user groups, the user can be granted the permissions of the roles.</li> <li>For example, add the user to the Hive user group to assign Hive permissions to the user.</li> </ul>                                                                                                                                                                                           |
| Primary Group    | Select a group as the primary group for the user to create directories and files. The drop-down list contains all groups selected in <b>User Group</b> .                                                                                                                                                                                                                                                                                                                                                                                                      |

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Role        | Click <b>Add</b> to bind a tenant role to the user.<br><b>NOTE</b> <ul style="list-style-type: none"><li>If a user wants to use the resources of tenant <b>tenant1</b> and to add or delete sub-tenants for <b>tenant1</b>, the user must be bound to both the <b>Manager_tenant</b> and <b>tenant1_Cluster ID</b> roles.</li><li>If the tenant has been associated with the HBase service and Ranger authentication is enabled for the cluster, you need to configure the HBase execution permissions on the Ranger page.</li></ul> |
| Description | Indicates the description of the current user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

**Step 3** Click **OK**.

----End

### 6.9.3.4 Adding an MRS Tenant Resource Pool

In an MRS cluster, users can logically divide Yarn cluster nodes to combine multiple NodeManagers into a Yarn resource pool. Each NodeManager belongs to one resource pool only. The system contains a **default** resource pool by default. All NodeManagers that are not added to customized resource pools belong to this resource pool.

You can create a customized resource pool on MRS and add hosts that have not been added to other customized resource pools to it.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Adding a Resource Pool on the MRS Management Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS details page, click **Tenant**.

**Step 4** Click the **Resource Pools** tab.

**Step 5** Click **Create Resource Pool**.

**Step 6** In **Create Resource Pool**, set the properties of the resource pool.

- Name:** Enter a name for the resource pool. The name of the newly created resource pool cannot be **default**. The name can contain digits, letters, and underscores (\_), and cannot start with an underscore (\_).

- **Resource Label:** Enter the resource label of the resource pool. The value can contain 1 to 50 characters, including digits, letters, underscores (\_), and hyphens (-), and must start with a digit or letter.
- **Available Hosts:** In the host list on the left, select a specified host name and add it to the resource pool. Only hosts in the cluster can be selected. The host list of a resource pool can be left blank.

**Step 7** Click **OK**.

**Step 8** After a resource pool is created, users can view the **Name**, **Members**, **Type**, **vCore** and **Memory** in the resource pool list. Hosts that are added to the customized resource pool are no longer members of the **default** resource pool.

----End

## Adding a Resource Pool on Manager

**Step 1** Log in to FusionInsight Manager. Choose **Tenant Resources**.

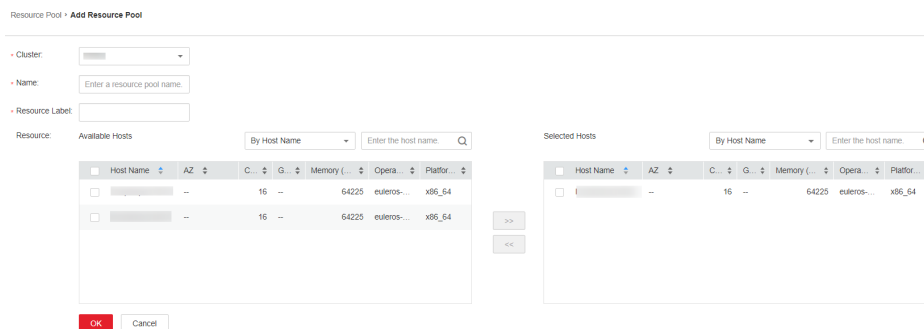
For MRS 2.x or earlier versions, select **Tenant**.

**Step 2** Click **Resource Pool** in the navigation pane, and click **Add Resource Pool**.

**Step 3** Set resource pool properties.

- **MRS cluster 3.x and later versions:**

**Figure 6-59** Adding resources



- **Cluster:** Select the cluster to which the resource pool is to be added.
- **Name:** Enter a name for the resource pool. The value can contain digits, letters, and underscores (\_), and cannot start with an underscore (\_).
- **Resource Label:** Enter the resource label of the resource pool. The value can contain 1 to 50 characters, including digits, letters, underscores (\_), and hyphens (-), and must start with a digit or letter.
- **Resource:** In the **Available Hosts** area, select specified hosts and click >> to add the hosts to the **Selected Hosts** area. Only hosts in the cluster can be selected. The host list of a resource pool can be left blank.

### NOTE

You can search for hosts by host name, number of CPU cores, memory, operating system, or platform type based on service requirements.

- **MRS 2.x and earlier versions:**



- **Name:** Enter a name for the resource pool. The name of the newly created resource pool cannot be **Default**.  
The name consists of 1 to 20 characters and can contain digits, letters, and underscores (\_) but cannot start with an underscore (\_).
- **Hosts:** In the host list on the left, select the name of a specified host and click >> to add the selected host to the resource pool. Only hosts in the cluster can be selected. The host list of a resource pool can be left blank.

**Step 4** Click **OK**.

After the resource pool is created, you can view it including the name, members, and mode in the resource pool list. Hosts that are added to the customized resource pool are no longer members of the **default** resource pool.

----End

### 6.9.3.5 Configuring the Queue Capacity Policy of a Resource Pool

After a resource pool is added, the capacity policies of available resources need to be configured for Yarn task queues. This ensures that tasks in the resource pool are running properly. Each queue can be configured with the queue capacity policy of only one resource pool. Users can view the queues in any resource pool and configure queue capacity policies. After the queue policies are configured, Yarn task queues and resource pools are associated.

You can configure queue policies on MRS.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).
- A resource pool has been added.
- The task queue is not associated with other resource pools except the default **default** resource pool.

### Configuring a Queue Capacity Policy on the MRS Management Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS details page, click **Tenant**.

**Step 4** Click the **Resource Distribution Policies** tab.

**Step 5** In **Resource Pools**, select a specified resource pool.

**Available Resource Quota:** indicates that all resources in each resource pool are available for queues by default.

**Step 6** Locate the specified queue in the **Resource Allocation** table, and click **Modify** in the **Operation** column.

**Step 7** In **Modify Resource Allocation**, configure the resource capacity policy of the task queue in the resource pool.

- **Capacity (%)**: specifies the percentage of the current tenant's computing resource usage.
- **Maximum Capacity (%)**: specifies the percentage of the current tenant's maximum computing resource usage.

**Step 8** Click **OK** to save the settings.

----End

## Configuring the Queue Capacity Policy on Manager

For clusters of MRS 3.x and later

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Tenant Resources > Dynamic Resource Plan** and click the **Resource Distribution Policy** tab.

**Step 3** Select the name of the target cluster from **Cluster** and select a resource pool from **Resource Pool**.

**Step 4** Locate the row that contains the target queue in the **Resource Allocation** area, and click **Modify** in the **Operation** column.

**Step 5** Modify the resource allocation.

- **Capacity scheduler**
  - **Capacity (%)**: indicates the percentage of compute resources used by the current tenant.
  - **Maximum Capacity (%)**: indicates the maximum percentage of computing resources used by the current tenant.
- **Superior Scheduler**
  - a. On the **Resource Configuration Policy** tab of the **Modify Resource Allocation** window, set the resource configuration policy of the queue in the resource pool.

**Figure 6-60** Resource configuration policy

Modify Resource Allocation ✕

Resource Configuration Policy    User Policy

---

\* Weight:

\* Minimum Resource:  %  (MB)  vCores  
Resources guaranteed for the tenant (preemption supported). The value can be a percentage of the parent tenant's resources or an absolute value. When a tenant has a light workload, the resources of the tenant are automatically allocated to other tenants. When the available resources of the tenant do not meet the minimum threshold, the tenant can preempt the resources lent to other tenants.

\* Maximum Resource:  %  (MB)  vCores  
Maximum resources that a tenant can use. The value can be a percentage of the parent tenant's resources or an absolute value.

\* Reserved Resource:  %  (MB)  vCores  
Reserved resources for the tenant. Even if there is no job in a tenant, the reserved resources cannot be used by other tenants. The value can be the percentage or absolute value of the parent tenant's resources.  
Note: The default resource pool capacity equals to set weight and minimum resource percentage value in advanced configuration. When a percentage and an absolute value are configured at the same time: vCore = max. (percentage, absolute value) and memory = max. (percentage, absolute value)

- **Weight:** The task queue with a larger weight preempts resources first when resources are insufficient. Its initial value is the same as the minimum resource percentage.
  - **Minimum Resource:** indicates the minimum resources that a tenant can obtain.
  - **Maximum Resource:** indicates the maximum resources that a tenant can obtain.
  - **Reserved Resource:** indicates the resources that are reserved for the tenant's queues and cannot be lent to other tenants' queues.
- b. In the **Modify Resource Allocation** dialog box, click the **User Policy** tab and set the user policy.

 **NOTE**

**defaultUser(built-in)** indicates that the policy specified by **defaultUser** is used if a user does not specify a policy. The default policy cannot be deleted.

- Click **Add User Policy** to add a user policy.
  - **Username:** indicates the name of a user.
  - **Weight:** The task queue with a larger weight preempts resources first when resources are insufficient.
  - **Max vCores:** indicates the maximum number of virtual cores that the user can obtain.
  - **Max Memory(MB):** indicates the maximum memory that the user can obtain.

- Click **Modify** in the **Operation** column to modify an existing user policy.
- Click **Clear** in the **Operation** column to delete an existing user policy.

**Step 6** Click **OK** to save the settings.

----End

**For MRS clusters of 2.x and earlier:**

**Step 1** On MRS Manager, click **Tenant**.

**Step 2** Click the **Dynamic Resource Plan** tab.

**Step 3** In **Resource Pools**, select a specified resource pool.

**Available Resource Quota:** indicates that all resources in each resource pool are available for queues by default.

**Step 4** Locate the specified queue in the **Resource Allocation** table, and click **Modify** in the **Operation** column.

**Step 5** In **Modify Resource Allocation**, configure the resource capacity policy of the task queue in the resource pool.

- **Capacity (%)**: specifies the percentage of the current tenant's computing resource usage.
- **Maximum Capacity (%)**: specifies the percentage of the current tenant's maximum computing resource usage.

**Step 6** Click **OK** to save the settings.

----End

### 6.9.3.6 Configuring the MRS Tenant Queue

You can modify the queue configuration of a specified tenant on MRS to scale in or out the queue based on service requirements. YARN queues are associated with resource pools for resource allocation and scheduling.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).
- A tenant associated with Yarn and allocated dynamic resources has been added.

### Configuring a Tenant Queue on the MRS Management Console

**Step 1** Log in to the MRS console.


**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS details page, click **Tenant**.

**Step 4** Click the **Queue Configuration** tab.

**Step 5** In the tenant queue table, click **Modify** in the **Operation** column of the specified tenant queue.

 **NOTE**

- In the tenant list on the left of the **Tenant** tab, click the target tenant. In the window that is displayed, choose **Resource**. On the page that is displayed, click  to open the queue modification page.
- A queue can be bound to only one non-default resource pool.
- **For MRS 2.x and earlier versions:**

**Table 6-50** Queue configuration parameters

| Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum Applications           | Specifies the maximum number of applications.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Maximum AM Resource Percent    | Specifies the maximum percentage of resources that can be used to run the ApplicationMaster in a cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Minimum User Limit Percent (%) | <p>Specifies the minimum percentage of resources consumed by a user.</p> <p>The resources for each user in a queue are limited at any time. If applications of multiple users are running at the same time in a queue, the resource usage of each user fluctuates between the minimum value and the maximum value. The minimum value is determined by the number of running applications, while the maximum value is determined by this parameter.</p> <p>For example, assume that this parameter is set to <b>25</b>. If two users submit applications to the queue, each user can use a maximum of 50% resources; if three users submit applications to the queue, each user can use a maximum of 33% resources; if four users submit applications to the queue, each user can use a maximum of 25% resources.</p> |
| User Limit Factor              | Specifies the limit factor of the maximum user resource usage. The maximum user resource usage percentage can be obtained by multiplying the limit factor with the percentage of the tenant's actual resource usage in the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Status                         | Specifies the current status of a resource plan.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| Parameter                                             | Description                                                                                                                                                                                                                                                 |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Default Resource Pool (Default Node Label Expression) | Specifies the resource pool used by a queue. The default value is <b>default</b> . If you want to change the resource pool, configure the queue capacity first. For details, see <a href="#">Configuring the Queue Capacity Policy of a Resource Pool</a> . |

- **For MRS cluster 3.x and later versions:**

**Table 6-51** Queue configuration parameters

| Parameter                 | Description                                                                                                                                                                                                                                                                                                         |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Max Master Shares (%)     | Indicates the maximum percentage of resources occupied by all ApplicationMasters in the current queue.                                                                                                                                                                                                              |
| Max Allocated vCores      | Indicates the maximum number of cores that can be allocated to a single YARN container in the current queue. The default value is <b>-1</b> , indicating that the number of cores is not limited within the value range.                                                                                            |
| Max Allocated Memory (MB) | Indicates the maximum memory that can be allocated to a single Yarn container in the current queue. The default value is <b>-1</b> , indicating that the memory is not limited within the value range.                                                                                                              |
| Max Running Apps          | Indicates the maximum number of tasks that can be executed at the same time in the current queue. The default value is <b>-1</b> , indicating that the number is not limited within the value range (the meaning is the same if the value is empty). The value <b>0</b> indicates that the task cannot be executed. |
| Max Running Apps per User | Maximum number of tasks that can be executed by each user in the current queue at the same time. The default value is <b>-1</b> , indicating that the number is not limited within the value range. If the value is <b>0</b> , the task cannot be executed.                                                         |
| Max Pending Apps          | Indicates the maximum number of tasks that can be suspended at the same time in the current queue. The default value is <b>-1</b> , indicating that the number is not limited within the value range (the meaning is the same if the value is empty). The value <b>0</b> indicates that tasks cannot be suspended.  |

| Parameter                | Description                                                                                                                                                                                                                                                                                                         |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Resource Allocation Rule | Indicates the rule for allocating resources to different tasks of a user. The rule can be FIFO or FAIR.<br><br>If a user submits multiple tasks in the current queue and the rule is FIFO, the tasks are executed one by one in sequential order. If the rule is FAIR, resources are evenly allocated to all tasks. |
| Default Resource Label   | Indicates that tasks are executed on a node with a specified resource label.<br><b>NOTE</b><br>If you need to use a new resource pool, change the default label to the new resource pool label.                                                                                                                     |
| Cross-Pool Scheduling    | Indicates whether containers in the current queue support cross-pool scheduling. This parameter is available for clusters of MRS 3.3.0 or later.<br><br>This function cannot be enabled for the default queue.                                                                                                      |
| Cross-Pool AM Scheduling | Indicates whether ApplicationMasters in the current queue support cross-pool scheduling. This parameter is available for clusters of MRS 3.3.0 or later.<br><br>This function cannot be enabled for the default queue.                                                                                              |
| Active                   | <ul style="list-style-type: none"><li>- <b>ACTIVE</b>: indicates that the current queue can receive and execute tasks.</li><li>- <b>INACTIVE</b>: indicates that the current queue can receive but cannot execute tasks. Tasks submitted to the queue are suspended.</li></ul>                                      |
| Open                     | <ul style="list-style-type: none"><li>- <b>OPEN</b>: indicates that the current queue is opened.</li><li>- <b>CLOSED</b>: indicates that the current queue is closed. Tasks submitted to the queue are rejected.</li></ul>                                                                                          |

----End

## Configuring a Tenant Queue on Manager

For clusters of MRS 3.x and later


**Step 1** On FusionInsight Manager, choose **Tenant Resources**.

**Step 2** Click the **Dynamic Resource Plan** tab.

**Step 3** Click the **Queue Configuration** tab.

**Step 4** Locate the row containing the specified tenant resource name and click **Modify** in the **Operation** column. Modify the parameters based on the scheduler type in use.

 **NOTE**

- YARN in a new cluster uses the Superior scheduler by default. You can switch the scheduler by referring to [Switching the MRS Tenant Resource Scheduler](#).
- To query the scheduler type, log in to Manager and search for the `yarn.resourcemanager.scheduler.class` parameter on the **All Configurations** page of the YARN service.
- You can also access the **Modify Queue Configuration** page as follows: In the tenant list on the **Tenant Resources Management** page, click the target tenant, click the **Resource** tab, and click  next to **Queue Configurations (Queue name)**.
- A queue can be bound to only one non-default resource pool. That is, a newly added resource pool can be bound to only one queue to serve as the default resource pool of the queue.
- For parameters such as **Max Allocated vCores**, **Max Allocated Memory(MB)**, **Max Running Apps**, **Max Running Apps per User**, and **Max Pending Apps**, if the value of a sub-tenant is **-1**, the value of the parent tenant can be set to a specific limit. If the parent tenant value is a specific limit, the sub-tenant value can be set to **-1**.
- **Max Allocated vCores** and **Max Allocated Memory(MB)** must be both changed to values other than **-1**.
- For queues with cross-resource-pool scheduling enabled, existing resource pools cannot be deleted during job running. Otherwise, running jobs may be continuously blocked because they cannot obtain resources. Similarly, if a new resource pool is configured for a queue during job running, the queue in the running state may not immediately use the resources in the new resource pool. The new resources are available only for jobs submitted after modification.
- To use a Superior scheduler cluster, configure sub-tenant properties by referring to the following table.

**Table 6-52** Queue configuration parameters

| Parameter                 | Description                                                                                                                                                                                                              |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Max Master Shares (%)     | Indicates the maximum percentage of resources occupied by all ApplicationMasters in the current queue.                                                                                                                   |
| Max Allocated vCores      | Indicates the maximum number of cores that can be allocated to a single YARN container in the current queue. The default value is <b>-1</b> , indicating that the number of cores is not limited within the value range. |
| Max Allocated Memory (MB) | Indicates the maximum memory that can be allocated to a single Yarn container in the current queue. The default value is <b>-1</b> , indicating that the memory is not limited within the value range.                   |



| Parameter                 | Description                                                                                                                                                                                                                                                                                                         |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Max Running Apps          | Indicates the maximum number of tasks that can be executed at the same time in the current queue. The default value is <b>-1</b> , indicating that the number is not limited within the value range (the meaning is the same if the value is empty). The value <b>0</b> indicates that the task cannot be executed. |
| Max Running Apps per User | Maximum number of tasks that can be executed by each user in the current queue at the same time. The default value is <b>-1</b> , indicating that the number is not limited within the value range. If the value is <b>0</b> , the task cannot be executed.                                                         |
| Max Pending Apps          | Indicates the maximum number of tasks that can be suspended at the same time in the current queue. The default value is <b>-1</b> , indicating that the number is not limited within the value range (the meaning is the same if the value is empty). The value <b>0</b> indicates that tasks cannot be suspended.  |
| Resource Allocation Rule  | Indicates the rule for allocating resources to different tasks of a user. The rule can be FIFO or FAIR.<br><br>If a user submits multiple tasks in the current queue and the rule is FIFO, the tasks are executed one by one in sequential order. If the rule is FAIR, resources are evenly allocated to all tasks. |
| Default Resource Label    | Indicates that tasks are executed on a node with a specified resource label.<br><br><b>NOTE</b><br>If you need to use a new resource pool, change the default label to the new resource pool label.                                                                                                                 |
| Cross-Pool Scheduling     | Indicates whether containers in the current queue support cross-pool scheduling. This parameter is available for clusters of MRS 3.3.0 or later.<br><br>This function cannot be enabled for the default queue.                                                                                                      |
| Cross-Pool AM Scheduling  | Indicates whether ApplicationMasters in the current queue support cross-pool scheduling. This parameter is available for clusters of MRS 3.3.0 or later.<br><br>This function cannot be enabled for the default queue.                                                                                              |

| Parameter | Description                                                                                                                                                                                                                                                                       |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Active    | <ul style="list-style-type: none"> <li>- <b>ACTIVE</b>: indicates that the current queue can receive and execute tasks.</li> <li>- <b>INACTIVE</b>: indicates that the current queue can receive but cannot execute tasks. Tasks submitted to the queue are suspended.</li> </ul> |
| Open      | <ul style="list-style-type: none"> <li>- <b>OPEN</b>: indicates that the current queue is opened.</li> <li>- <b>CLOSED</b>: indicates that the current queue is closed. Tasks submitted to the queue are rejected.</li> </ul>                                                     |

- For a cluster using the Capacity scheduler, **Tenant Resource Name (Queue)** indicates the tenant and queue name. Configure sub-tenants properties by referring to the following table.

**Table 6-53** Queue configuration parameters

| Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum Applications           | Specifies the maximum number of applications.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Maximum AM Resource Percent    | Specifies the maximum percentage of resources that can be used to run the ApplicationMaster in a cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Minimum User Limit Percent (%) | <p>Specifies the minimum percentage of resources consumed by a user.</p> <p>The resources for each user in a queue are limited at any time. If applications of multiple users are running at the same time in a queue, the resource usage of each user fluctuates between the minimum value and the maximum value. The minimum value is determined by the number of running applications, while the maximum value is determined by this parameter.</p> <p>For example, assume that this parameter is set to <b>25</b>. If two users submit applications to the queue, each user can use a maximum of 50% resources; if three users submit applications to the queue, each user can use a maximum of 33% resources; if four users submit applications to the queue, each user can use a maximum of 25% resources.</p> |
| User Limit Factor              | Specifies the limit factor of the maximum user resource usage. The maximum user resource usage percentage can be obtained by multiplying the limit factor with the percentage of the tenant's actual resource usage in the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Parameter                                             | Description                                                                                                                                                                                                                                                 |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status                                                | Specifies the current status of a resource plan.                                                                                                                                                                                                            |
| Default Resource Pool (Default Node Label Expression) | Specifies the resource pool used by a queue. The default value is <b>default</b> . If you want to change the resource pool, configure the queue capacity first. For details, see <a href="#">Configuring the Queue Capacity Policy of a Resource Pool</a> . |

**Step 5** Click **OK**.

----End

**For MRS clusters of 2.x and earlier:**

**Step 1** On MRS Manager, click **Tenant**.

**Step 2** Click the **Dynamic Resource Plan** tab.

**Step 3** Click the **Queue Configuration** tab.

**Step 4** In the tenant queue table, click **Modify** in the **Operation** column of the specified tenant queue.

 **NOTE**

In the tenant list on the left of the **Tenant** tab, click the target tenant. In the window that is displayed, choose **Resource**. On the page that is displayed, click the edit icon to open the queue modification page.

**Table 6-54** Queue configuration parameters

| Parameter                   | Description                                                                                               |
|-----------------------------|-----------------------------------------------------------------------------------------------------------|
| Maximum Applications        | Specifies the maximum number of applications.                                                             |
| Maximum AM Resource Percent | Specifies the maximum percentage of resources that can be used to run the ApplicationMaster in a cluster. |

| Parameter                                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Minimum User Limit Percent (%)                        | <p>Specifies the minimum percentage of resources consumed by a user.</p> <p>The resources for each user in a queue are limited at any time. If applications of multiple users are running at the same time in a queue, the resource usage of each user fluctuates between the minimum value and the maximum value. The minimum value is determined by the number of running applications, while the maximum value is determined by this parameter.</p> <p>For example, assume that this parameter is set to <b>25</b>. If two users submit applications to the queue, each user can use a maximum of 50% resources; if three users submit applications to the queue, each user can use a maximum of 33% resources; if four users submit applications to the queue, each user can use a maximum of 25% resources.</p> |
| User Limit Factor                                     | <p>Specifies the limit factor of the maximum user resource usage. The maximum user resource usage percentage can be obtained by multiplying the limit factor with the percentage of the tenant's actual resource usage in the cluster.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Status                                                | <p>Specifies the current status of a resource plan.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Default Resource Pool (Default Node Label Expression) | <p>Specifies the resource pool used by a queue. The default value is <b>default</b>. If you want to change the resource pool, configure the queue capacity first. For details, see <a href="#">Configuring the Queue Capacity Policy of a Resource Pool</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

----End

## 6.9.4 Managing MRS Tenant Resources

### 6.9.4.1 Managing the MRS Tenant Resource Directory

You can manage the HDFS storage directory used by a specific tenant on MRS. The management operations include adding a tenant directory, modifying the directory file quota, modifying the storage space, and deleting a directory.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

- A tenant associated with HDFS storage resources has been added.

## Managing a Tenant Directory

**Step 1** Go to the **Tenant** page.

- On the MRS management console: Click the target cluster. On the cluster details page, click **Tenant**.
- On the Manager of an MRS 3.x or later cluster: Click **Tenant Resources**.
- On the Manager of an MRS 2.x or later cluster: Click **Tenant**.

**Step 2** In the tenant list on the left, click the target tenant.

**Step 3** Click the resource tab and perform the following operations to manage tenant directory:

- **View the tenant directory.**

View related information in the **HDFS Storage** pane.

- The **File Number Threshold** column provides the quota for files and directories of the tenant directory. (On the MRS management console, this parameter is **Maximum Number of Files/Directories**.)
- The **Storage Space Quota** column indicates storage space size of tenant directories.

- **Add a tenant directory.**

a. In the **HDFS Storage** pane, click **Create Directory**.

- **Parent Directory:** Select a storage directory of a parent tenant.

This parameter is not displayed if the current tenant is not a sub-tenant. If the parent tenant has multiple directories, select any of them.

- **Path:** Enter the tenant directory.

- If the current tenant is not a sub-tenant, the new path is created in the HDFS root directory.
- If the current tenant is a sub-tenant, the new path is created in the specified directory.

A complete HDFS storage directory can contain a maximum of 1,023 characters. An HDFS directory name contains digits, letters, spaces, and underscores (\_). The name cannot start or end with a space.

- **Maximum Number of Files/Directories:** Enter the quota of files and directories. This parameter is optional.
- **File Number Threshold (%)** is valid only when **Quota** is set. If the ratio of the number of used files to the value of **Quota** exceeds the value of this parameter, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario. This parameter is available for clusters of MRS 3.x or later.

### NOTE

The number of used files is collected every hour. Therefore, the alarm indicating that the ratio of used files exceeds the threshold is delayed.

- **Storage Space Quota:** Specify the storage space size of the tenant directory.

To ensure data reliability, one copy of a file is automatically generated when the file is stored in HDFS. That is, two copies of the same file are stored by default. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value of **Storage Space Quota** is set to **500**, the actual space for storing files is about 250 MB ( $500/2 = 250$ ).

- **Storage Space Threshold (%):** If the ratio of used storage space to the value of **Space Quota** exceeds this value, an alarm is generated. If this parameter is not specified, no alarm will be reported in this scenario. This parameter is available for clusters of MRS 3.x or later.

 **NOTE**

The used storage space is collected every hour. Therefore, the alarm indicating that the ratio of used storage space exceeds the threshold is delayed.

- b. Click **OK**. The system creates tenant directories in the HDFS root directory.
- **Modify a tenant directory.**
    - a. In the **HDFS Storage** table, click **Modify** in the **Operation** column of the specified tenant directory.
    - b. Change the parameter values based on service requirements.
    - c. Click **OK**.
  - **Delete a tenant directory.**
    - a. In the **HDFS Storage** table, click **Delete** in the **Operation** column of the specified tenant directory.

The default HDFS storage directory set during tenant creation cannot be deleted. Only the newly added HDFS storage directory can be deleted.
    - b. Click **OK**. The tenant directory is deleted.

----End

### 6.9.4.2 Managing MRS Tenant Resource Pools

When hosts in a resource pool need to be adjusted based on service requirements, you can modify members in the resource pool on MRS.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).
- To delete a resource pool, the following requirements must be met:
  - Any queue in a cluster cannot use the resource pool to be deleted as the default resource pool. Before deleting the resource pool, cancel the

default resource pool. For details, see [Configuring the MRS Tenant Queue](#).

- Resource distribution policies of all queues have been cleared from the resource pool being deleted. For details, see [Clearing the MRS Tenant Queue Configuration](#).

## Managing Resource Pools on the MRS Management Console

**Step 1** Log in to the MRS console.


**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the MRS details page, click **Tenant**.

**Step 4** Click the **Resource Pools** tab.

**Step 5** Locate the row that contains the specified resource pool, and click **Modify** in the **Operation** column.

**Step 6** Manage hosts in a resource pool.

- Adding a host: In the host list on the left, select the specified host name and add it to the resource pool.
- Deleting a host: In the host list on the right, click  next to a host to remove the host from the resource pool. The host list of a resource pool can be left blank.

**Step 7** Click **OK**.

**Step 8** Locate the row that contains the target resource pool in the **Resource Pools** tab, and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.

----End

## Managing Resource Pools on Manager

**Step 1** Log in to Manager. Choose **Tenant Resources > Resource Pool**.

For MRS 2.x or earlier versions, choose **Tenant > Resource Pool**.

**Step 2** Locate the row that contains the specified resource pool, and click **Edit** in the **Operation** column.

For MRS 2.x and earlier versions and MRS 3.3.0 and later versions, click **Modify**.

**Step 3** Manage hosts in a resource pool.

- Adding a host: Select the name of a specified host in host list on the left and click >> to add the selected host to the resource pool.
- Deleting a host: In the host list on the right, select the name of a specified host and click << to add the selected host to the resource pool. The host list of a resource pool can be left blank.

**Step 4** Click **OK**.

**Step 5** Locate the row that contains the target resource pool in the **Resource Pools** tab, and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.

----End

### 6.9.4.3 Clearing the MRS Tenant Queue Configuration

Users can clear the configuration of a queue on MRS Manager when the queue does not need resources from a resource pool or if a resource pool needs to be disassociated from the queue. Clearing queue configurations means that the resource capacity policy of the queue is canceled.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).
- If a queue is to be unbound from a resource pool, this resource pool cannot serve as the default resource pool of the queue. Therefore, you must first change the default resource pool of the queue to another one. For details, see [Configuring the MRS Tenant Queue](#).

#### Clearing Queue Configuration

**Step 1** Go to the resource distribution policy page.

- On the Manager, the page path is as follows:
  - For MRS 3.x and later versions, choose **Tenant Resources > Dynamic Resource Plan > Resource Distribution Policy**.
  - For MRS 2.x and earlier versions, choose **Tenant > Dynamic Resource Plan > Resource Distribution Policy**.
- On the MRS console: Go to the cluster details page, click **Tenants** and then **Resource Distribution Policies**.

**Step 2** In the **Resource Pools** tab, select a specified resource pool.

**Step 3** Locate the specified queue in the **Resource Allocation** table, and click **Clear** in the **Operation** column.

In the displayed dialog box, click **Yes** to clear the queue configurations from the current resource pool.

#### NOTE

If no resource capacity policy is configured for a queue, the clearing function is unavailable for the queue by default.

----End



### 6.9.4.4 Restoring MRS Tenant Data After YARN Is Reinstalled

Tenant data is stored on Manager and in cluster components by default. When components are restored from faults or reinstalled, some tenant configuration data may be abnormal. In this case, you can manually restore the tenant data.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

#### Restoring Tenant Data

**Step 1** Go to the **Tenant** page.


- On the MRS management console: Click the target cluster. On the cluster details page, click **Tenant**.
- On the Manager of an MRS 3.x or later cluster: Click **Tenant Resources**.
- On the Manager of an MRS 2.x or later cluster: Click **Tenant**.

**Step 2** In the tenant list on the left, click a tenant node.

**Step 3** Check the status of the tenant data.

1. In the **Summary** tab, check **Tenant Status**. A green icon indicates that the tenant is available and gray indicates that the tenant is unavailable.
2. Click **Resources** and check the status of **Yarn** or **HDFS Storage**. Green indicates that the resource is available, and gray indicates that the resource is unavailable.
3. Click **Service Association** and check the **Status** column of the associated service table. **Good** indicates that the component can provide services for the associated tenant. **Bad** indicates that the component cannot provide services for the tenant.
4. If any check result is abnormal, go to **Step 4** to restore tenant data.

**Step 4** Click **Restore Tenant Data**.

To restore data on the Manager of an MRS 3.x or later cluster, click . In the displayed dialog box, enter the password of the current user and click **OK**.

**Step 5** In the displayed window, select one or more components whose data needs to be restored. Click **OK**. The system automatically restores the tenant data.

----End

### 6.9.4.5 Deleting an MRS Tenant

You can delete MRS tenants that are no longer used based on service requirements, releasing resources occupied by the tenants.

## Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).
- You have checked whether the tenant to be deleted has sub-tenants. If the tenant has sub-tenants, delete them; otherwise, you cannot delete the tenant.
- The role of the tenant to be deleted cannot be associated with any user or user group.

## Deleting a Tenant on the MRS Management Console

**Step 1** On the MRS details page, click **Tenants**.

**Step 2** In the tenant list on the left, move the cursor to the tenant node to be deleted and click **Delete**.

The **Delete Tenant** dialog box is displayed. If you want to save the tenant data, select **Reserve the data of this tenant**. Otherwise, the tenant's storage space will be deleted.

**Step 3** Click **OK**.

It takes a few minutes to save the configuration. After the tenant is deleted successfully, the role and storage space of the tenant are also deleted.

### NOTE


- After the tenant is deleted, the task queue of the tenant still exists in YARN.
- If you choose not to reserve data when deleting the parent tenant, data of sub-tenants is also deleted if the sub-tenants use storage resources.

----End

## Deleting a Tenant on Manager

**For MRS 3.x and later versions:**

**Step 1** Log in to Manager and choose **Tenant Resources**.

**Step 2** In the tenant list on the left, click the target tenant and click .

### NOTE

- If you want to retain the tenant data, select **Reserve the data of this tenant resource**. Otherwise, the storage space of the tenant will be deleted.
- To delete a tenant without retaining the tenant data as a user who does not belong to the supergroup, you should first log in to the HDFS client as a user who belongs to the supergroup and then manually clear the storage space of that tenant to avoid residual data.

**Step 3** In the **Delete Tenant** dialog box, enter **DELETE** in the confirmation text box. Click **OK**.

It takes a few minutes to save the configuration. After the tenant is deleted successfully, the role and storage space of the tenant are also deleted.

 NOTE

After the tenant is deleted, the task queue of the tenant still exists in YARN. The queue of the tenant is not displayed on the role management page in YARN.

----End

**For MRS 2.x and earlier versions:**

**Step 1** On MRS Manager, click **Tenant**.

**Step 2** In the tenant list on the left, move the cursor to the tenant node to be deleted and click **Delete**.

The **Delete Tenant** dialog box is displayed. If you want to save the tenant data, select **Reserve the data of this tenant**. Otherwise, the tenant's storage space will be deleted.

**Step 3** Click **OK**.

It takes a few minutes to save the configuration. After the tenant is deleted successfully, the role and storage space of the tenant are also deleted.

 NOTE

- After the tenant is deleted, the task queue of the tenant still exists in YARN.
- If you choose not to reserve data when deleting the parent tenant, data of sub-tenants is also deleted if the sub-tenants use storage resources.

----End

## 6.9.4.6 Managing Global User Policies When Using Superior Scheduler

### Scenario

If a tenant uses a Superior scheduler, you can configure the global policy for users to use the resource scheduler, including:

- Maximum running apps
- Maximum pending apps
- Default queue

### Procedure

- Add a policy.
  - a. On FusionInsight Manager, choose **Tenant Resources**.
  - b. Choose **Dynamic Resource Plan**.
  - c. Click the **Global User Policy** tab.

 NOTE

**defaults(default setting)** indicates that the policy specified for **defaults** is used if a user does not have a global policy. The default policy cannot be deleted.

- d. Click **Create Global User Policy**. In the displayed dialog box, set the following parameters:

**Figure 6-61** Creating a global user policy

### Global User Policy

\* Cluster:

\* Username:

Max Running Apps:

Max Pending Apps:

Default Queue:

- **Cluster:** Select the target cluster.
  - **Username:** indicates the user for whom resource scheduling is controlled. Enter an existing username in the current cluster.
  - **Max Running Apps:** indicates the maximum number of tasks that the user can run in the current cluster.
  - **Max Pending Apps:** indicates the maximum number of tasks that the user can suspend in the current cluster.
  - **Default Queue:** indicates the queue of the user. Enter the name of an existing queue in the current cluster.
- Modify a policy.
    - a. On FusionInsight Manager, choose **Tenant Resources**.
    - b. Choose **Dynamic Resource Plan**.
    - c. Click the **Global User Policy** tab.
    - d. In the row that contains the desired user policy, click **Modify** in the **Operation** column.
    - e. In the displayed dialog box, modify parameters and click **OK**.
  - Delete a policy.
    - a. On FusionInsight Manager, choose **Tenant Resources**.
    - b. Choose **Dynamic Resource Plan**.
    - c. Click the **Global User Policy** tab.
    - d. In the row that contains the desired user policy, click **Delete** in the **Operation** column.  
In the displayed dialog box, click **OK**.

## 6.9.4.7 Clearing Tenant's Non-Associated Queues Using Capacity Scheduler

### Scenario

If Yarn uses the Capacity scheduler, deleting a tenant only sets the queue capacity of the tenant to **0** and the tenant status to **STOPPED** but does not clear the queues of the tenant in Yarn. Limited by the Yarn mechanism, queues cannot be dynamically deleted. You can run commands to manually delete residual queues.

### Impact on the System

- During the script execution, the Controller service is restarted, Yarn configurations are synchronized, and the active and standby ResourceManagers are restarted.
- FusionInsight Manager becomes inaccessible during the restart of the Controller service.
- After the active and standby ResourceManagers are restarted, an alarm is generated indicating that Yarn and components that depend on Yarn are temporarily unavailable.

### Prerequisites

Queues of a deleted tenant still exist.

### Procedure

**Step 1** Check that queues of the deleted tenant still exist.

1. On FusionInsight Manager, choose **Cluster**, click the name of the target cluster, and choose **Services > Yarn**. Click the link of the active ResourceManager in **ResourceManager WebUI** to go to the ResourceManager web UI.
2. Click **Scheduler** in the navigation tree on the left. In the right pane, you can view that queues of the tenant still exist in the **STOPPED** state and their **Configured Capacity** is **0**.

**Step 2** Log in to the active management node as user **omm**.

**Step 3** Switch the directory and execute the **cleanQueuesAndRestartRM.sh** script.

```
cd ${BIGDATA_HOME}/om-server/om/sbin
./cleanQueuesAndRestartRM.sh -c Cluster ID
```

#### NOTE

You can choose **Cluster**, click the cluster name, and choose **Cluster Properties** on FusionInsight Manager to view the cluster ID.

During the script execution, you need to enter **yes** and the password.

```
Running the script will restart Controller and restart ResourceManager.
Are you sure you want to continue connecting (yes/no)?yes
Please input admin password:
Begin to backup queues ...
...
```

- Step 4** After the script is executed successfully, log in to FusionInsight Manager, choose **Cluster**, click the cluster name, and choose **Services > Yarn**. Click the link of the active ResourceManager in **ResourceManager WebUI** to go to the ResourceManager web UI.
- Step 5** Click **Scheduler** in the navigation tree on the left. In the right pane, you can view that queues of the tenant have been cleared.
- End

## 6.9.5 Switching the MRS Tenant Resource Scheduler

### Scenario

The newly installed MRS cluster uses the Superior scheduler by default. If the cluster is upgraded from an earlier version, you can switch the YARN scheduler from the Capacity scheduler to the Superior scheduler with a few clicks.

### Prerequisites

- The network connectivity of the cluster is proper and secure, and the YARN service status is normal.
- During scheduler switching, tenants cannot be added, deleted, or modified. In addition, services cannot be started or stopped.

### Switching from the Capacity Scheduler to the Superior Scheduler

#### Impact on the system

- Because the ResourceManager is restarted during scheduler switching, submitting jobs to YARN will fail at that time.
- During scheduler switching, tasks in a job being executed on YARN will continue, but new tasks cannot be started.
- After scheduler switching is complete, jobs executed on YARN may fail, causing service interruptions.
- After scheduler switching is complete, parameters of the Superior scheduler are used for tenant management.
- After scheduler switching is complete, tenant queues whose capacity is 0 in the Capacity scheduler cannot be allocated resources in the Superior scheduler. As a result, jobs submitted to these tenant queues fail to be executed. Therefore, you are advised not to set the capacity of a tenant queue to 0 in the Capacity scheduler.
- After scheduler switching is complete, you cannot add or delete resource pools, YARN node labels, or tenants during the observation period. If such an operation is performed, the scheduler cannot be rolled back to the Capacity scheduler.

#### NOTE

The recommended observation period for scheduler switching is one week. If resource pools, YARN node labels, or tenants are added or deleted during this period, the observation period ends immediately.

- Rollback may cause the loss of partial or all Yarn job information.

### Procedure

**Step 1** Modify YARN service parameters and ensure that the YARN service status is normal.

1. Log in to FusionInsight Manager as an administrator.
2. Log in to FusionInsight Manager and choose **Cluster > Services > Yarn**. Click **Configurations** then **All Configurations**, search for **yarn.resourcemanager.webapp.pagination.enable**, and check whether the value is **true**.
  - If yes, go to **Step 1.3**.
  - If no, set the parameter to **true** and click **Save** to save the configuration. On the **Dashboard** tab page of YARN, choose **More > Restart Service**, verify the identity, and click **OK**. After the service is restarted, go to **Step 1.3**.
3. Choose **Cluster > Name of the desired cluster > Services**, and check whether the YARN service status is normal.

**Step 2** Log in to the active management node as user **omm**.

**Step 3** Switch the scheduler.

The following switching modes are available:

**0**: converts the Capacity scheduler configurations into the Superior scheduler configurations and then switches the Capacity scheduler to the Superior scheduler.

**1**: converts the Capacity scheduler configurations into the Superior scheduler configurations only.

**2**: switches the Capacity scheduler to the Superior scheduler only.

- Mode **0** is recommended if the cluster environment is simple and the number of tenants is less than 20.

Run the following command:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/switchScheduler.sh -c Cluster ID -m 0
```

#### NOTE

You can choose **Cluster**, click the cluster name, and choose **Cluster Properties** on FusionInsight Manager to view the cluster ID.

```
Start to convert Capacity scheduler to Superior Scheduler, clusterId=1
Start to convert Capacity scheduler configurations to Superior. Please wait...
Convert configurations successfully.
Start to switch the Yarn scheduler to Superior. Please wait...
Switch the Yarn scheduler to Superior successfully.
```

- If the cluster environment or tenant information is complex and you need to retain the queue configurations of the Capacity scheduler on the Superior scheduler, it is recommended that you use mode **1** first to convert the Capacity scheduler configurations, check the converted configurations, and then use mode **2** to switch the Capacity scheduler to the Superior scheduler.
  - a. Run the following command to convert the Capacity scheduler configurations into the Superior scheduler configurations:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/switchScheduler.sh -c
Cluster ID -m 1
```

```
Start to convert Capacity scheduler to Superior Scheduler, clusterId=1
Start to convert Capacity scheduler configurations to Superior. Please wait...
Convert configurations successfully.
```

- b. Run the following command to switch the Capacity scheduler to the Superior scheduler:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/switchScheduler.sh -c
Cluster ID -m 2
```

```
Start to convert Capacity scheduler to Superior Scheduler, clusterId=1
Start to switch the Yarn scheduler to Superior. Please wait...
Switch the Yarn scheduler to Superior successfully.
```

- If you do not need to retain the queue configurations of the Capacity scheduler, use mode 2.
  - a. Log in to FusionInsight Manager and delete all tenants except the default tenant.
  - b. On FusionInsight Manager, delete all resource pools except the default resource pool.

Run the following command to switch the Capacity scheduler to the Superior scheduler:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/switchScheduler.sh -c
Cluster ID -m 2
```

```
Start to convert Capacity scheduler to Superior Scheduler, clusterId=1
Start to switch the Yarn scheduler to Superior. Please wait...
Switch the Yarn scheduler to Superior successfully.
```

#### NOTE

You can query the scheduler switching logs on the active management node.

- `${BIGDATA_LOG_HOME}/controller/aos/switch_scheduler.log`
- `${BIGDATA_LOG_HOME}/controller/aos/aos.log`

----End

## Rollback Operations

You can manually switch the Superior scheduler back to the Capacity scheduler. However, this operation is only a workaround and is not allowed in most cases.

If the customer has special requirements for switching back to the Capacity scheduler, the following conditions must be met:

- The observation period has not expired.
- No resource pool, YARN node label, or tenant is added or deleted during the observation period.

---

#### NOTICE

If resource pools, YARN node labels, or tenants are added or deleted, resource pools or queues may not exist after the Superior scheduler is switched back to the Capacity scheduler. As a result, the Capacity scheduler cannot run properly.

---



The procedure is as follows:

**Step 1** Change the scheduler to the Capacity scheduler and start YARN.

1. Log in to FusionInsight Manager.
2. Go to the **Configurations** page of YARN and modify the parameters listed in [Table 6-55](#).

**Table 6-55** Modifying YARN configuration items

| Parameter                                     | Description                                                                        |
|-----------------------------------------------|------------------------------------------------------------------------------------|
| yarn.resourcemanager.scheduler.class          | org.apache.hadoop.yarn.server.resourcemanager.scheduler.capacity.CapacityScheduler |
| yarn.http.rmwebapp.external.classes           | Left empty                                                                         |
| hadoop.http.rmwebapp.scheduler.page.classes   | Left empty                                                                         |
| yarn.resourcemanager.webapp.pagination.enable | false                                                                              |

3. Click **Save** and then click **OK** in the displayed dialog box.
4. Roll restart the YARN service, enter the password, and click **OK**.

**Step 2** Log in to the active management node and restart the AOS service.

1. Log in to the active OMS server as user **omm** using PuTTY.
2. Run the following command to disable logout upon timeout:

```
TMOUT=0
```

 **NOTE**

After the operations in this section are complete, run the **TMOUT=Timeout interval** command to restore the timeout interval in a timely manner. For example, **TMOUT=600** indicates that a user is logged out if the user does not perform any operation within 600 seconds.

3. Run the following command to restart the AOS service:  
**`\${BIGDATA\_HOME}/om-server/om/sbin/aos\_cmd.sh restart**

----End

## 6.10 Managing MRS Cluster Users

### 6.10.1 Cluster User Permissions

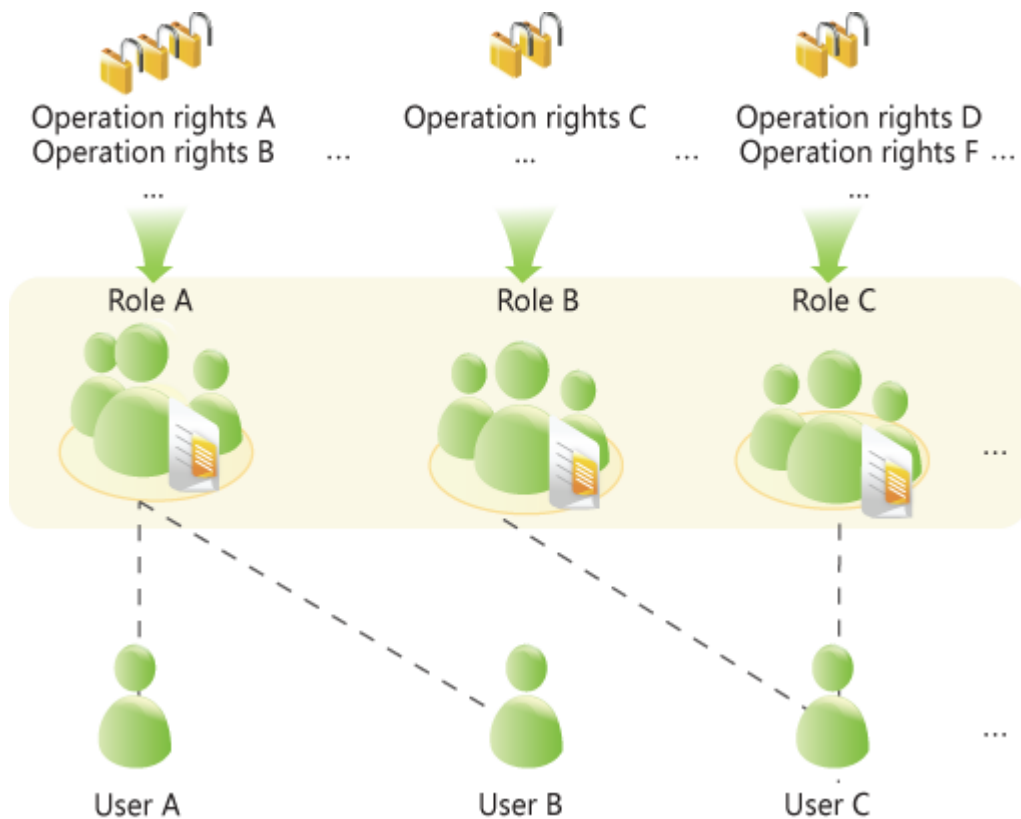
### 6.10.1.1 MRS Cluster User Permission Model

#### Role-based Access Control

MRS adopts the role-based access control (RBAC) mode to manage rights on the big data system. It integrates the right management functions of the components to centrally manage rights. Common users are shielded from internal right management details, and the right management operations are simplified for administrators, improving right management usability and user experience.

The right model of MRS consists four parts: users, user groups, roles, and rights.

**Figure 6-62** Right model



- **Right**

Right, which is defined by components, allows users to access a certain resource of one component. Different components have different rights for their resources.

For example:

- HDFS provides read, write, and execute permissions on files.
- HBase provides create, read, and write permissions on tables.

- **Role**

Role is a collection of component rights. Each role can have multiple rights of multiple components. Different roles can have the rights of a resource of one component.

- **User group**

User group is a collection of users. When a user group is bound to a role, users in this group obtain the rights defined by the role.

Different user groups can be associated with the same role. A user group can also be associated with no role, and this user group does not have the rights of any component resources.

 **NOTE**

In some components, the system grants related rights to specific user groups by default.

- **User**

A user is a visitor to the system. Each user has the rights of the user group and role associated with the user. Users need to be added to the user group or associated with roles to obtain the corresponding rights.

## Policy-based Access Control

The Ranger component uses policy-based access control (PBAC) to manage rights and implement fine-grained data access control on components such as HDFS, Hive, and HBase.

 **NOTE**

The component supports only one right control mechanism. After the Ranger right control policy is enabled for the component, the right on the component in the role created on FusionInsight Manager becomes invalid (The ACL rules of HDFS and Yarn still take effect). You need to add a policy on the Ranger management page to grant rights on resources.

The Ranger right model consists of multiple right policies. A right policy consists of the following parts:

- **Resource**

Resources are provided by components and can be accessed by users, such as HDFS files or folders, queues in Yarn, and databases, tables, and columns in Hive.

- **User**

A User is a visitor to the system. The rights of each user are obtained based on the policy associated with the user. Information about users, user groups, and roles in the LDAP is periodically synchronized to the Ranger.

- **Permission**

In a policy, you can configure various access conditions for resources, such as file read and write, permission conditions, rejection conditions, and exception conditions.

## Permission Mechanism

MRS adopts the Lightweight Directory Access Protocol (LDAP) to store data of users and user groups. Information about role definitions is stored in the relational database and the mapping between roles and rights is saved in components.

MRS uses Kerberos for unified authentication.

The user permission verification process is as follows:

1. A client (a user terminal or MRS component service) invokes the MRS authentication interface.
2. MRS uses the login username and password for Kerberos authentication.
3. If the authentication succeeds, the client sends a request for accessing the server (an MRS component service).
4. The server finds the user group and role to which the login user belongs.
5. The server obtains all rights of the user group and the role.
6. The server checks whether the client has the right to access the resources it applies for.

**Example (RBAC):**

There are three files in HDFS, that is, fileA, fileB, and fileC.

- roleA has read and write right for fileA, and roleB has the read right for fileB.
- groupA is bound to roleA, and groupB is bound to roleB.
- userA belongs to groupA and roleB, and userB belongs to groupB.

When userA successfully logs in to the system and accesses the HDFS:

1. HDFS obtains the role (roleB) to which userA is bound.
2. HDFS also obtains the role (roleA) to which the user group of userA is bound.
3. In this case, userA has all the rights of roleA and roleB.
4. As a result, userA has read and write rights for fileA, has the read right on fileB, and has no right for fileC.

Similarly, when userB successfully logs in to the system and accesses the HDFS:

1. userB only has the rights of roleB.
2. As a result, userB has the read right on fileB, and has no rights for fileA and fileC.

### 6.10.1.2 MRS Cluster User Identity Authentication Policy

The big data platform performs user identity authentication to prevent invalid users from accessing the cluster. The cluster provides authentication capabilities in both security mode and normal mode.

#### Security Mode

The clusters in security mode use the Kerberos authentication protocol for security authentication. The Kerberos protocol supports mutual authentication between clients and servers. This eliminates the risks incurred by sending user credentials over the network for simulated authentication. In clusters, KrbServer provides the Kerberos authentication support.

**Kerberos user object**

In the Kerberos protocol, each user object is a principal. A complete principal consists of username and domain name. In O&M or application development scenarios, the user identity must be verified before a client connects to a server. Users for O&M and service operations are classified into human-machine and machine-machine users. The password of human-machine users is manually

configured, while the password of machine-machine users is generated by the system randomly.

### Kerberos authentication

Kerberos supports password and keytab authentication. The validity period of authentication is 24 hours by default.

- Password authentication: User identity is verified by entering the correct password. This mode mainly used in O&M scenarios where human-machine users are used. The configuration command is **kinit** *Username*.
- Keytab authentication: Keytab files contain users' principal and encrypted credential information. When keytab files are used for authentication, the system automatically uses encrypted credential information to perform authentication and the user password does not need to be entered. This mode is mainly used in component application development scenarios where machine-machine users are used. Keytab authentication can also be configured using the **kinit** command.

## Normal Mode

Different components in a normal cluster use the native open-source authentication mode and do not support the **kinit** authentication command. FusionInsight Manager (including DBService, KrbServer, and LdapServer) uses the username and password for authentication. [Table 6-56](#) lists the authentication modes used by components.

**Table 6-56** Component authentication modes

| Service    | Authentication Mode                                                                                                 |
|------------|---------------------------------------------------------------------------------------------------------------------|
| IoTDB      | Simple authentication                                                                                               |
| CDL        | No authentication                                                                                                   |
| ClickHouse | Simple authentication                                                                                               |
| Flume      | No authentication                                                                                                   |
| HBase      | <ul style="list-style-type: none"><li>• Web UI: No authentication</li><li>• Client: simple authentication</li></ul> |
| HDFS       | <ul style="list-style-type: none"><li>• Web UI: no authentication</li><li>• Client: simple authentication</li></ul> |
| HetuEngine | <ul style="list-style-type: none"><li>• Web UI: no authentication</li><li>• Client: no authentication</li></ul>     |
| Hive       | Simple authentication                                                                                               |
| Hue        | Username and password authentication                                                                                |
| Kafka      | No authentication                                                                                                   |

| Service   | Authentication Mode                                                                                                                    |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------|
| Loader    | <ul style="list-style-type: none"><li>• Web UI: username and password authentication</li><li>• Client: no authentication</li></ul>     |
| MapReduce | <ul style="list-style-type: none"><li>• Web UI: no authentication</li><li>• Client: no authentication</li></ul>                        |
| Oozie     | <ul style="list-style-type: none"><li>• Web UI: username and password authentication</li><li>• Client: simple authentication</li></ul> |
| Spark2x   | <ul style="list-style-type: none"><li>• Web UI: no authentication</li><li>• Client: simple authentication</li></ul>                    |
| Storm     | No authentication                                                                                                                      |
| YARN      | <ul style="list-style-type: none"><li>• Web UI: no authentication</li><li>• Client: simple authentication</li></ul>                    |
| ZooKeeper | Simple authentication                                                                                                                  |

The authentication modes are as follows:

- Simple authentication: When the client connects to the server, the client automatically authenticates the user (for example, the OS user **root** or **omm**) by default. The authentication is imperceptible to the administrator or service user, which does not require **kinit**.
- Username and password authentication: Use the username and password of human-machine users in the cluster for authentication.
- No authentication: Any user can access the server by default.

### 6.10.1.3 MRS Cluster User Permission Authentication Policy

#### Security Mode

After a user is authenticated by the big data platform, the system determines whether to verify the user's permission based on the actual permission management configuration to ensure that the user has limited or all permission on resources. If the user does not have the permission for accessing cluster resources, the system administrator must grant the required permission to the user. Otherwise, the user fails to access the resources. The cluster provides permission verification capabilities in both security mode and normal mode. The specific permission items of the components are the same in the two modes.

By default, the Ranger service is installed and Ranger authentication is enabled for a newly installed cluster in security mode. You can set fine-grained security access policies for accessing component resources through the permission plug-in of the component. If Ranger authentication is not required, administrators can manually disable it on the service page. After Ranger authentication is disabled, the system continues to perform permission control based on the role model of FusionInsight Manager when accessing component resources.

In a cluster in security mode, the following components support Ranger authentication: HDFS, YARN, Kafka, Hive, HBase, Storm, Impala, HetuEngine, CDL, and Spark2x.

For a cluster upgraded from an earlier version, Ranger authentication is not used by default when users access component resources. The administrator can manually enable Ranger authentication after installing Ranger.

By default, all components in the cluster of the security edition authenticate access. The authentication function cannot be disabled.

## Normal Mode

Different components in a normal cluster use their own native open-source authentication behavior. [Table 6-57](#) lists detailed permission verification modes.

In a normal cluster, Ranger supports permission control on component resources based on OS users. The following components support Ranger authentication: HBase, HDFS, Hive, Spark2x, and YARN.

**Table 6-57** Component permission verification modes in normal clusters

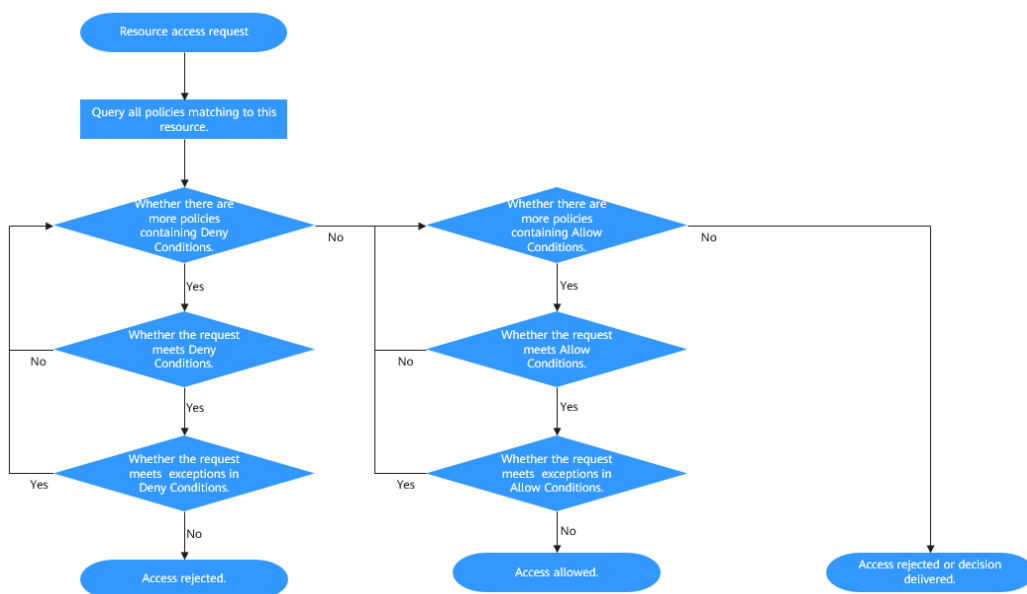
| Service    | Permission Verification | Permission Verification Enabling and Disabling |
|------------|-------------------------|------------------------------------------------|
| IoTDB      | Required                | Not supported                                  |
| ClickHouse | Required                | Not supported                                  |
| Flume      | Not required            | Not supported                                  |
| HBase      | Not required            | Supported                                      |
| HDFS       | Required                | Supported                                      |
| HetuEngine | Not required            | Not supported                                  |
| Hive       | Not required            | Not supported                                  |
| Hue        | Not required            | Not supported                                  |
| Kafka      | Not required            | Not supported                                  |
| Loader     | Not required            | Not supported                                  |
| MapReduce  | Not required            | Not supported                                  |
| Oozie      | Required                | Not supported                                  |
| Spark2x    | Not required            | Not supported                                  |
| Storm      | Not required            | Not supported                                  |
| YARN       | Not required            | Supported                                      |
| ZooKeeper  | Required                | Supported                                      |
| CDL        | Not required            | Not supported                                  |

## Condition Priorities of the Ranger Permission Policy

When configuring a permission policy for a resource, you can configure Allow Conditions, Exclude from Allow Conditions, Deny Conditions, and Exclude from Deny Conditions for the resource, to meet unexpected requirements in different scenarios.

The priorities of different conditions are listed in descending order: Exclude from Deny Conditions > Deny Conditions > Exclude from Allow Conditions > Allow Conditions

The following figure shows the process of determining condition priorities. If the component resource request does not match the permission policy in Ranger, the system rejects the access by default. However, for HDFS and Yarn, the system delivers the decision to the access control layer of the component for determination.



For example, if you want to grant the read and write permissions of the **FileA** folder to the **groupA** user group, but the user in the group is not **UserA**, you can add an allowed condition and an exception condition.

### 6.10.1.4 Default Permissions of the MRS Cluster

#### Role

| Default Role          | Description                                                                                                                                                                       |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Manager_administrator | Manager administrator who has all permissions for Manager.<br>Manager administrators can create first-level tenants, create and modify user groups, and specify user permissions. |
| Manager_operator      | Manager operator who has all the permissions on the <b>Homepage</b> , <b>Cluster</b> , <b>Hosts</b> , and <b>O&amp;M</b> tab pages.                                               |



| Default Role              | Description                                                                                                                                                                                                                                                                      |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Manager_auditor           | Manager auditor who has all permissions on the <b>Audit</b> tab page.<br>Manager auditors can view and manage Manager system audit logs.                                                                                                                                         |
| Manager_viewer            | Manager viewer who has the permission to view information about <b>Homepage, Cluster, Hosts, Alarm, Events</b> , and <b>System &gt; Permission</b> , and download clients. (Only MRS 3.2.0 or later supports client download.)                                                   |
| Manager_tenant            | Manager tenant administrator.<br>This role can create and manage sub-tenants for the non-leaf tenants to which the current user belongs. It has the permission to view alarms and events on <b>O&amp;M &gt; Alarm</b> .                                                          |
| System_administrator      | System administrator, this role has Manager system administrator rights and all services administrator rights.                                                                                                                                                                   |
| default                   | This role is the default role created for the <b>default</b> tenant. It has the management permissions on the Yarn component and the default queue. The default role of the default tenant that is not the first cluster to be installed is <b>c&lt;cluster ID&gt;_default</b> . |
| Manager_administrator_180 | FusionInsight Manager System administrator group. Internal system user group, which is used only between components.                                                                                                                                                             |
| Manager_auditor_181       | FusionInsight Manager system auditor group. Internal system user group, which is used only between components.                                                                                                                                                                   |
| Manager_operator_182      | FusionInsight Manager system operator group. Internal system user group, which is used only between components.                                                                                                                                                                  |
| Manager_viewer_183        | FusionInsight Manager system viewer group. Internal system user group, which is used only between components.                                                                                                                                                                    |
| System_administrator_186  | System administrator group. Internal system user group, which is used only between components.                                                                                                                                                                                   |
| Manager_tenant_187        | Tenant system user group. Internal system user group, which is used only between components.                                                                                                                                                                                     |
| default_1000              | This group is created for tenant. Internal system user group, which is used only between components.                                                                                                                                                                             |

## User group

| Type          | Default User Group                                                                                                                      | Description                                                                                                                                                                                                                                                                                                                                      |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OS User Group | hadoop                                                                                                                                  | Users added to this group are granted the permission to submit all Yarn queue tasks.                                                                                                                                                                                                                                                             |
|               | hadoopmanager                                                                                                                           | Users added to this user group can have the O&M manager rights of HDFS and Yarn. The O&M manager of HDFS can access the NameNode WebUI and perform active to standby switchover manually. The O&M manager of Yarn can access the ResourceManager WebUI, operate NodeManager nodes, refresh queues, and set node labels, but cannot submit tasks. |
|               | hetuadmin                                                                                                                               | HetuEngine administrator group. Users in this group have the permission to perform operations on HSConsole.                                                                                                                                                                                                                                      |
|               | hive                                                                                                                                    | Common user group. Hive users must belong to this user group.                                                                                                                                                                                                                                                                                    |
|               | iotdbgroup                                                                                                                              | Users added to this user group have the administrator rights of the IoTDB component.                                                                                                                                                                                                                                                             |
|               | kafka                                                                                                                                   | Kafka common user group. A user in this group can access a topic only when a user in the kafkaadmin group grants the read and write permission of the topic to the user.                                                                                                                                                                         |
|               | kafkaadmin                                                                                                                              | Kafka administrator group. Users in this group have the rights to create, delete, authorize, read, and write all topics.                                                                                                                                                                                                                         |
|               | kafkasuperuser                                                                                                                          | Topic read/write user group of Kafka. Users added to this group have the read and write permissions on all topics.                                                                                                                                                                                                                               |
|               | cdladmin                                                                                                                                | CDL administrator group. Only users in this group can access CDL APIs.                                                                                                                                                                                                                                                                           |
|               | cdl                                                                                                                                     | Common user group of CDL. Users in this group can create and query CDL jobs.                                                                                                                                                                                                                                                                     |
|               | storm                                                                                                                                   | Users who are added to the storm user group can submit topologies and manage their own topologies.                                                                                                                                                                                                                                               |
|               | stormadmin                                                                                                                              | Users who are added to the stormadmin user group can have the storm administrator rights and can submit topologies and manage all topologies.                                                                                                                                                                                                    |
|               | supergroup                                                                                                                              | Users added to this user group have the administrator rights of HBase, HDFS, and Yarn and can use Hive.                                                                                                                                                                                                                                          |
| yarnviewgroup | Indicates the read-only user group of the Yarn task. Users in this user group can have the view permission on Yarn and MapReduce tasks. |                                                                                                                                                                                                                                                                                                                                                  |

| Type          | Default User Group | Description                                                                                                                                                                                                                                         |
|---------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | check_sec_ldap     | Perform internal test on the active LDAP to see whether it works properly. This user group is generated randomly in a test and automatically deleted after the test is complete. Internal system user group, which is used only between components. |
|               | compcommon         | System internal group for accessing cluster system resources. All system users and system running users are added to this user group by default.                                                                                                    |
| OS User Group | wheel              | Primary group of the FusionInsight internal running user omm.                                                                                                                                                                                       |
|               | ficommon           | System common group that corresponds to <b>compcommon</b> for accessing cluster common resource files stored in the OS.                                                                                                                             |

 **NOTE**

If the current cluster is not the cluster that is installed for the first time in FusionInsight Manager, the default user group name of all components except Manager in the cluster is *c<cluster ID>\_ default user group name*, for example, **c2\_hadoop**.

## User

For details, see [MRS Cluster User Accounts](#).

## Service-related User Security Parameters

- **HDFS**  
The **dfs.permissions.superusergroup** parameter specifies the administrator group with the highest permission on the HDFS. The default value is **supergroup**.
- **Spark2x**  
The **spark.admin.acls** parameter specifies the administrator list of the Spark2x. Members in the list are authorized to manage all Spark tasks. Users not added in the list cannot manage all Spark tasks. The default value is **admin**.

### 6.10.1.5 Synchronizing IAM Users to MRS

IAM user synchronization is to synchronize IAM users bound with MRS policies to the MRS system and create accounts with the same usernames but different passwords as the IAM users. Then, you can use an IAM username (the password needs to be reset by user **admin** of Manager) to log in to Manager for cluster management, and submit jobs on the GUI in a cluster with Kerberos authentication enabled.

**Table 6-58** compares IAM users' permission policies and the synchronized users' permissions on MRS. For details about the default permissions on Manager, see [Default Permissions of the MRS Cluster](#).

**Table 6-58** Policy and permission mapping after synchronization

| Policy Type  | IAM Policy           | User's Default Permissions on MRS After Synchronization                                                                                                                                                                                                         | Have Permission to Perform the Synchronization | Have Permission to Submit Jobs |
|--------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|--------------------------------|
| Fine-grained | MRS ReadOnlyAccess   | Manager_viewer                                                                                                                                                                                                                                                  | No                                             | No                             |
|              | MRS CommonOperations | <ul style="list-style-type: none"> <li>• Manager_viewer</li> <li>• default</li> <li>• launcher-job</li> </ul>                                                                                                                                                   | No                                             | Yes                            |
|              | MRS FullAccess       | <ul style="list-style-type: none"> <li>• Manager_administrator</li> <li>• Manager_auditor</li> <li>• Manager_operator</li> <li>• Manager_tenant</li> <li>• Manager_viewer</li> <li>• System_administrator</li> <li>• default</li> <li>• launcher-job</li> </ul> | Yes                                            | Yes                            |

| Policy Type | IAM Policy                                                | User's Default Permissions on MRS After Synchronization                                                                                                                                                                                                         | Have Permission to Perform the Synchronization | Have Permission to Submit Jobs |
|-------------|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|--------------------------------|
| RBAC        | MRS Administrator                                         | <ul style="list-style-type: none"> <li>• Manager_administrator</li> <li>• Manager_auditor</li> <li>• Manager_operator</li> <li>• Manager_tenant</li> <li>• Manager_viewer</li> <li>• System_administrator</li> <li>• default</li> <li>• launcher-job</li> </ul> | No                                             | Yes                            |
|             | Server Administrator, Tenant Guest, and MRS Administrator | <ul style="list-style-type: none"> <li>• Manager_administrator</li> <li>• Manager_auditor</li> <li>• Manager_operator</li> <li>• Manager_tenant</li> <li>• Manager_viewer</li> <li>• System_administrator</li> <li>• default</li> <li>• launcher-job</li> </ul> | Yes                                            | Yes                            |

| Policy Type | IAM Policy           | User's Default Permissions on MRS After Synchronization                                                                                                                                                                                                         | Have Permission to Perform the Synchronization                                                                                                                                                                                                                                            | Have Permission to Submit Jobs |
|-------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
|             | Tenant Administrator | <ul style="list-style-type: none"> <li>• Manager_administrator</li> <li>• Manager_auditor</li> <li>• Manager_operator</li> <li>• Manager_tenant</li> <li>• Manager_viewer</li> <li>• System_administrator</li> <li>• default</li> <li>• launcher-job</li> </ul> | Yes                                                                                                                                                                                                                                                                                       | Yes                            |
| Custom      | Custom policy        | <ul style="list-style-type: none"> <li>• Manager_viewer</li> <li>• default</li> <li>• launcher-job</li> </ul>                                                                                                                                                   | <ul style="list-style-type: none"> <li>• If custom policies use RBAC policies as a template, refer to the RBAC policies.</li> <li>• If custom policies use fine-grained policies as a template, refer to the fine-grained policies. The fine-grained policies are recommended.</li> </ul> | Yes                            |

 NOTE

To facilitate user permission management, use fine-grained policies rather than RBAC policies. In fine-grained policies, the Deny action takes precedence over other actions.

- A user has permission to synchronize IAM users only when the user has the Tenant Administrator role or has the Server Administrator, Tenant Guest, and MRS Administrator roles at the same time.
- A user with the **action:mrs:cluster:syncUser** policy has permission to synchronize IAM users.

## Synchronizing IAM Users

- Step 1** Create a user and authorize the user to use MRS. For details, see [Creating an IAM User and Granting MRS Permissions](#).
- Step 2** Log in to the MRS management console and create a cluster. For details, see [Manually Buying an MRS Cluster](#).
- Step 3** In the navigation pane on the left, choose **Active Clusters**. Click a cluster name to go to the cluster details page.
- Step 4** On the **Dashboard** tab page, click **Synchronize** next to **IAM User Sync** to synchronize IAM users.
- Step 5** In the **IAM User Sync** dialog box, search for the user group to which the IAM user to be synchronized belongs and click the user group name. In the **User** column, select the desired IAM user and click **Synchronize**.

 NOTE

- You can select all users to synchronize them at a time.
  - If you select user groups only, users will not be synchronized. You must select specific user names in the user group.
  - All user groups are displayed. Those cannot be selected cannot be synchronized.
- Step 6** After a synchronization request is sent, choose **Operation Logs** in the navigation tree on the left of the MRS console to check whether the synchronization is successful. For details about the logs, see [Viewing MRS Operation Logs](#).
  - Step 7** After the synchronization is successful, use the user synchronized with IAM to perform subsequent operations.

 NOTE

- When the policy of the user group to which the IAM user belongs changes from **MRS ReadOnlyAccess** to **MRS CommonOperations**, **MRS FullAccess**, or **MRS Administrator**, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** (System Security Services Daemon) cache of cluster nodes needs time to be updated. Then, submit a job. Otherwise, the job may fail to be submitted.
- When the policy of the user group to which the IAM user belongs changes from **MRS CommonOperations**, **MRS FullAccess**, or **MRS Administrator** to **MRS ReadOnlyAccess**, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** cache of cluster nodes needs time to be updated.
- After you click **Synchronize** on the right side of **IAM User Sync**, the cluster details page is blank for a short time, because user data is being synchronized. The page will be properly displayed after the data synchronization is complete.

- Submitting jobs in a security cluster: Users can submit jobs using the job management function on the GUI in the security cluster. For details, see [Running a MapReduce Job](#).
- All tabs are displayed on the cluster details page, including **Components**, **Tenants**, and **Backups & Restorations**.
- Logging in to Manager
  - a. Log in to Manager as user **admin**. For details, see [Accessing MRS Manager](#).
  - b. Initialize the password of the user synchronized with IAM. For details, see [Initializing MRS Cluster User Passwords](#).
  - c. Modify the role bound to the user group to which the user belongs to control user permissions on Manager. For details, see [Managing MRS Cluster User Groups](#). For details about how to create and modify a role, see [Creating a Role](#). After the component role bound to the user group to which the user belongs is modified, it takes some time for the role permissions to take effect.
  - d. Log in to Manager using the user synchronized with IAM and the password after the initialization in [Step 7.b](#).

 **NOTE**

If the IAM user's permission changes, go to [Step 4](#) to perform second synchronization. After the second synchronization, a system user's permissions are the union of the permissions defined in the IAM system policy and the permissions of roles added by the system user on Manager. After the second synchronization, a custom user's permissions are subject to the permissions configured on Manager.

- System user: If all user groups to which an IAM user belongs are bound to system policies (RABC policies and fine-grained policies belong to system policies), the IAM user is a system user.
- Custom user: If the user group to which an IAM user belongs is bound to any custom policy, the IAM user is a custom user.

#### **Step 8** Undo IAM user synchronization.

To undo the synchronization of an IAM user, select the user in the **User** column in the **Synchronized** tab and click **Undo Sync**.

To undo the synchronization of all users in an IAM user group, select the user group in the **User Group** column in the **Synchronized** tab and click **Undo Sync**.

----End

## 6.10.2 MRS Cluster User Accounts

This section describes information about default users in MRS clusters.

### Account List (MRS 3.x and Later Versions)

- **User types**

The MRS cluster provides the following three types of users. The system administrator needs to periodically change the passwords. It is not recommended to use the default passwords.



| User Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| System users          | <ul style="list-style-type: none"> <li>● User created on FusionInsight Manager for O&amp;M and service scenarios. There are two types of users:                             <ul style="list-style-type: none"> <li>– <b>Human-machine</b> user: used in scenarios such as FusionInsight Manager O&amp;M and operations on a component client. When creating a user of this type, you need to set password and confirm password by referring to <a href="#">Creating an MRS Cluster User</a>.</li> <li>– <b>Machine-machine</b> user: used for system application development.</li> </ul> </li> <li>● User who runs OMS processes</li> </ul> |
| Internal system users | Internal user to perform Kerberos authentication, process communications, save user group information, and associate user permissions. It is recommended that internal system users not be used in O&M scenarios. Operations can be performed as user <b>admin</b> or another user created by the system administrator based on service requirements.                                                                                                                                                                                                                                                                                       |
| Database users        | <ul style="list-style-type: none"> <li>● User who manages OMS database and accesses data</li> <li>● User who runs service components (Hue, Hive, HetuEngine, Loader, Oozie, Ranger, JobGateway, and DBService) in the database.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                  |

- **System user**

 **NOTE**

- User **root** of the OS is required, the password of user **root** on all nodes must be the same.
- User **ldap** of the OS is required. Do not delete this account. Otherwise, the cluster may not work properly. The OS administrator maintains the password management policies.

| User Type            | Username | Initial Password      | Description                                                                                                                                                                                                                                                                                                                                     | Password Change Method                                                                                 |
|----------------------|----------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| System administrator | admin    | User-defined password | <p>FusionInsight Manager administrator.</p> <p><b>NOTE</b><br/>By default, user <b>admin</b> does not have the management permission on other components. For example, when accessing the native UI of a component, the user fails to access the complete component information due to insufficient management permission on the component.</p> | For details, see <a href="#">Changing or Resetting the Password for User admin of an MRS Cluster</a> . |
| Node OS user         | ommdba   | Random password       | User that creates the system database. This user is an OS user generated on the management node and does not require a unified password. This account cannot be used for remote login.                                                                                                                                                          | For details, see <a href="#">Changing the Passwords for OS Users of an MRS Cluster Node</a> .          |
|                      | omm      | Random password       | Internal running user of the system. This user is an OS user generated on all nodes and does not require a unified password.                                                                                                                                                                                                                    |                                                                                                        |

- **Internal system user**

| User Type                  | Default User             | Initial Password                                                                                                                                                  | Description                                                                  | Password Change Method                                                                                                                                                                                                                                                    |
|----------------------------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Kerberos administrator     | kadmin/admin             | Admin@123                                                                                                                                                         | Used to add, delete, modify, and query user accounts on Kerberos.            | For details, see <a href="#">Changing the Password for the Kerberos Administrator of an MRS Cluster.</a>                                                                                                                                                                  |
| OMS Kerberos administrator | kadmin/admin             | Admin@123                                                                                                                                                         | Used to add, delete, modify, and query user accounts on OMS Kerberos.        | For details, see <a href="#">Changing the Password of the OMS Kerberos Administrator.</a>                                                                                                                                                                                 |
| LDAP administrator         | cn=root,dc=hadoop,dc=com | <ul style="list-style-type: none"> <li>Versions earlier than MRS 3.1.2: LdapChangeMe@123</li> <li>MRS 3.1.2 or later: randomly generated by the system</li> </ul> | Used to add, delete, modify, and query the user account information on LDAP. | <ul style="list-style-type: none"> <li>For versions earlier than MRS 3.1.2, see <a href="#">Changing the Password for a Regular LDAP User of an MRS Cluster.</a></li> <li>For MRS 3.1.2 or later, see <a href="#">Modifying the OMS Service Configuration.</a></li> </ul> |

| User Type              | Default User                              | Initial Password                                                                                                                                                  | Description                                                                      | Password Change Method |
|------------------------|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|------------------------|
| OMS LDAP administrator | cn=root,dc=hadoop,dc=com                  | <ul style="list-style-type: none"> <li>Versions earlier than MRS 3.1.2: LdapChangeMe@123</li> <li>MRS 3.1.2 or later: randomly generated by the system</li> </ul> | Used to add, delete, modify, and query the user account information on OMS LDAP. |                        |
| LDAP user              | cn=pg_search_dn,ou=Users,dc=hadoop,dc=com | Randomly generated by the system                                                                                                                                  | Used to query information about users and user groups on LDAP.                   |                        |
| OMS LDAP user          | cn=pg_search_dn,ou=Users,dc=hadoop,dc=com | Randomly generated by the system                                                                                                                                  | Used to query information about users and user groups on OMS LDAP.               |                        |

| User Type                  | Default User                        | Initial Password                                                                                                                                                  | Description                                                          | Password Change Method                                                                                                                                                                                                                                                |
|----------------------------|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LDAP administrator account | cn=krbkdc,ou=Users,dc=hadoop,dc=com | <ul style="list-style-type: none"> <li>Versions earlier than MRS 3.1.2: LdapChangeMe@123</li> <li>MRS 3.1.2 or later: randomly generated by the system</li> </ul> | Used to query Kerberos component authentication account information. | <ul style="list-style-type: none"> <li>For versions earlier than MRS 3.1.2, see <a href="#">Changing the LDAP Administrator Password for an MRS Cluster</a>.</li> <li>For MRS 3.1.2 or later, see <a href="#">Modifying the OMS Service Configuration</a>.</li> </ul> |

| User Type | Default User                         | Initial Password                                                                                                                                                       | Description                                                                                   | Password Change Method |
|-----------|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|------------------------|
|           | cn=krbadmin,ou=Users,dc=hadop,dc=com | <ul style="list-style-type: none"> <li>• Versions earlier than MRS 3.1.2: Ldap ChangeMe@123</li> <li>• MRS 3.1.2 or later: randomly generated by the system</li> </ul> | Used to add, delete, modify, and query Kerberos component authentication account information. |                        |

| User Type              | Default User | Initial Password | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Password Change Method                                                                                  |
|------------------------|--------------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| Component running user | hdfs         | Hdfs@123         | <p>This user is the HDFS system administrator and has the following permissions:</p> <ol style="list-style-type: none"> <li>1. File system operation permissions: <ul style="list-style-type: none"> <li>• Views, modifies, and creates files.</li> <li>• Views and creates directories.</li> <li>• Views and modifies the groups where files belong.</li> <li>• Views and sets disk quotas for users.</li> </ul> </li> <li>2. HDFS management operation permissions: <ul style="list-style-type: none"> <li>• Views the web UI status.</li> <li>• Views and sets the active and standby HDFS status.</li> <li>• Enters and exits the HDFS in security mode.</li> <li>• Checks the HDFS file system.</li> </ul> </li> <li>3. Logs in to the FTP service page.</li> </ol> | <p>For details, see <a href="#">Changing the Passwords for MRS Cluster Component Running Users</a>.</p> |

| User Type | Default User | Initial Password | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Password Change Method |
|-----------|--------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | hbase        | Hbase@123        | <p>This user is the HBase and HBase1 to HBase4 system administrator and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Cluster management permission: Performs <b>Enable</b> and <b>Disable</b> operations on tables to trigger MajorCompact and ACL operations.</li> <li>• Grants and revokes permissions, and shuts down the cluster.</li> <li>• Table management permission: Creates, modifies, and deletes tables.</li> <li>• Data management permission: Reads data in tables, column families, and columns.</li> <li>• Logs in to the HMaster web UI.</li> <li>• Logs in to the FTP service page.</li> </ul> |                        |



| User Type | Default User | Initial Password | Description                                                                          | Password Change Method |
|-----------|--------------|------------------|--------------------------------------------------------------------------------------|------------------------|
|           | cdl          | CDCUser123!      | System administrator of the CDL<br>Currently, CDL does not involve user permissions. |                        |

| User Type | Default User | Initial Password | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Password Change Method |
|-----------|--------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | iotdb        | lotdb@123        | <p>This user is the IoTDB system administrator and has the following user permissions:</p> <ol style="list-style-type: none"> <li>1. IoTDB administrator permissions: <ul style="list-style-type: none"> <li>• Creates or deletes a storage group.</li> <li>• Uses TTL.</li> </ul> </li> <li>2. IoTDB data operation permissions: <ul style="list-style-type: none"> <li>• Creates, modifies, and deletes a time sequence.</li> <li>• Writes, reads, and deletes data in a time sequence.</li> </ul> </li> <li>3. Views user or role permission information.</li> <li>4. Grants or revokes permissions to or from a user or role.</li> </ol> |                        |

| User Type | Default User | Initial Password | Description                                                                                                                                                                                                                                                                                                                            | Password Change Method |
|-----------|--------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           |              |                  | <p><b>NOTE</b><br/>In a normal cluster, the IoTDB service retains the features of open-source versions. The default username is <b>root</b>, and the default password is <b>root</b>. This user is an administrator and has all permissions, which cannot be assigned, revoked, or deleted.</p>                                        |                        |
|           | mapred       | Mapred@123       | <p>This user is the MapReduce system administrator and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Submits, stops, and views the MapReduce tasks.</li> <li>• Modifies the Yarn configuration parameters.</li> <li>• Logs in to the FTP service page.</li> <li>• Logs in to the Yarn web UI.</li> </ul> |                        |

| User Type | Default User  | Initial Password  | Description                                                                                                                                                                                                                                                                  | Password Change Method |
|-----------|---------------|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | zookeeper     | ZooKeeper@123     | <p>This user is the ZooKeeper system administrator and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Adds, deletes, modifies, and queries all nodes in ZooKeeper.</li> <li>• Modifies and queries quotas of all nodes in ZooKeeper.</li> </ul> |                        |
|           | rangeradmin   | Rangeradmin@123   | <p>This user has the Ranger system management permissions and user permissions:</p> <ul style="list-style-type: none"> <li>• Ranger web UI management permission</li> <li>• Management permission of each component that uses Ranger authentication</li> </ul>               |                        |
|           | rangerauditor | Rangerauditor@123 | Default audit user of the Ranger system.                                                                                                                                                                                                                                     |                        |

| User Type | Default User | Initial Password | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Password Change Method |
|-----------|--------------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | hive         | Hive@123         | <p>This user is the Hive system administrator and has the following permissions:</p> <ol style="list-style-type: none"> <li>1. Hive administrator permissions: <ul style="list-style-type: none"> <li>• Creates, deletes, and modifies a database.</li> <li>• Creates, queries, modifies, and deletes a table.</li> <li>• Queries, inserts, and uploads data.</li> </ul> </li> <li>2. HDFS file operation permissions: <ul style="list-style-type: none"> <li>• Views, modifies, and creates files.</li> <li>• Views and creates directories.</li> <li>• Views and modifies the groups where files belong.</li> </ul> </li> <li>3. Submits and stops the MapReduce tasks.</li> <li>4. Ranger policy management permission</li> </ol> |                        |

| User Type | Default User   | Initial Password                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Password Change Method |
|-----------|----------------|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | kafka          | Kafka@123                        | <p>This user is the Kafka system administrator and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Creates, deletes, produces, and consumes the topic; modifies the topic configuration.</li> <li>• Controls the cluster metadata, modifies the configuration, migrates the replica, elects the leader, and manages ACL.</li> <li>• Submits, queries, and deletes the consumer group offset.</li> <li>• Queries the delegation token.</li> <li>• Queries and submits the transaction.</li> </ul> |                        |
|           | storm          | Admin@123                        | <p>Storm system administrator</p> <p>User permission: Submits Storm tasks.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                               |                        |
|           | rangerusersync | Randomly generated by the system | Synchronizes users and internal users of user groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                        |

| User Type | Default User                            | Initial Password                 | Description                                                                                                                                                                                                                                                                                                                                                                                                            | Password Change Method |
|-----------|-----------------------------------------|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | rangertag sync                          | Randomly generated by the system | Internal user for synchronizing tags.                                                                                                                                                                                                                                                                                                                                                                                  |                        |
|           | rangerobs / hadoop.<System domain name> | Randomly generated by the system | System administrator used by Guardian to access Ranger                                                                                                                                                                                                                                                                                                                                                                 |                        |
|           | jobserver                               | Randomly generated by the system | JobGateway system administrator, who has the following permissions: <ol style="list-style-type: none"> <li>HDFS file operation permissions:                             <ul style="list-style-type: none"> <li>Views, modifies, and creates files.</li> <li>Views and creates directories.</li> <li>Views and modifies the groups where files belong.</li> </ul> </li> <li>Manager administrator permission</li> </ol> |                        |
|           | HTTP/_HOST                              | Randomly generated by the system | Internal user of the JobGateway service, which is used for Kerberos authentication of the HTTP service                                                                                                                                                                                                                                                                                                                 |                        |
|           | oms/manager                             | Randomly generated by the system | Controller and NodeAgent authentication user. The user has the permission on the <b>supergroup</b> group.                                                                                                                                                                                                                                                                                                              |                        |

| User Type | Default User       | Initial Password                 | Description                                                                                                                                                                                                                                                                                 | Password Change Method |
|-----------|--------------------|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | backup/<br>manager | Randomly generated by the system | User for running backup and restoration tasks. The user has the permission on the <b>supergroup</b> , <b>wheel</b> , and <b>ficommon</b> groups. After cross-system mutual trust is configured, the user has the permission to access data in the HDFS, HBase, Hive, and ZooKeeper systems. |                        |



| User Type | Default User                              | Initial Password                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Password Change Method |
|-----------|-------------------------------------------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | hdfs/hadoop.< <i>System domain name</i> > | Randomly generated by the system | <p>This user is used to start the HDFS and has the following permissions:</p> <ol style="list-style-type: none"> <li>1. File system operation permissions: <ul style="list-style-type: none"> <li>• Views, modifies, and creates files.</li> <li>• Views and creates directories.</li> <li>• Views and modifies the groups where files belong.</li> <li>• Views and sets disk quotas for users.</li> </ul> </li> <li>2. HDFS management operation permissions: <ul style="list-style-type: none"> <li>• Views the web UI status.</li> <li>• Views and sets the active and standby HDFS status.</li> <li>• Enters and exits the HDFS in security mode.</li> <li>• Checks the HDFS file system.</li> </ul> </li> <li>3. Logs in to the FTP service page.</li> </ol> |                        |

| User Type | Default User                            | Initial Password                 | Description                                                                                                                                                                                                                                                                                                                     | Password Change Method |
|-----------|-----------------------------------------|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | hetuser/<br>hadoop.<System domain name> | Randomly generated by the system | <p>This user is used to start HetuEngine and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Accesses KrbServer and HDFS files in the cluster from HetuEngine.</li> <li>• Used for communication between HetuEngine internal nodes.</li> </ul>                                                      |                        |
|           | mapred/<br>hadoop.<System domain name>  | Randomly generated by the system | <p>This user is used to start the MapReduce and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Submits, stops, and views the MapReduce tasks.</li> <li>• Modifies the Yarn configuration parameters.</li> <li>• Logs in to the FTP service page.</li> <li>• Logs in to the Yarn web UI.</li> </ul> |                        |
|           | mr_zk/<br>hadoop.<System domain name>   | Randomly generated by the system | Used for MapReduce to access ZooKeeper.                                                                                                                                                                                                                                                                                         |                        |

| User Type | Default User                        | Initial Password                 | Description                                                                                                                                                                                                                                                                   | Password Change Method |
|-----------|-------------------------------------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | hbase/hadoop.<System domain name>   | Randomly generated by the system | User for the authentication between internal components during the HBase system startup.                                                                                                                                                                                      |                        |
|           | hbase/zkclient.<System domain name> | Randomly generated by the system | User for HBase to perform ZooKeeper authentication in a security mode cluster.                                                                                                                                                                                                |                        |
|           | thrift/hadoop.<System domain name>  | Randomly generated by the system | ThriftServer system startup user.                                                                                                                                                                                                                                             |                        |
|           | thrift/<hostname>                   | Randomly generated by the system | User for the ThriftServer system to access HBase. This user has the read, write, execution, creation, and administration permission on all NameSpaces and tables of HBase. <hostname> indicates the name of the host where the ThriftServer node is installed in the cluster. |                        |

| User Type | Default User                     | Initial Password                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Password Change Method |
|-----------|----------------------------------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | hive/hadoop.<system domain name> | Randomly generated by the system | <p>User for the authentication between internal components during the Hive system startup. The user permissions are as follows:</p> <ol style="list-style-type: none"> <li>1. Hive administrator permissions: <ul style="list-style-type: none"> <li>• Creates, deletes, and modifies a database.</li> <li>• Creates, queries, modifies, and deletes a table.</li> <li>• Queries, inserts, and uploads data.</li> </ul> </li> <li>2. HDFS file operation permissions: <ul style="list-style-type: none"> <li>• Views, modifies, and creates files.</li> <li>• Views and creates directories.</li> <li>• Views and modifies the groups where files belong.</li> </ul> </li> <li>3. Submits and stops the MapReduce tasks.</li> </ol> |                        |

| User Type | Default User                                | Initial Password                 | Description                                                                                                                            | Password Change Method |
|-----------|---------------------------------------------|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | loader/hadoop.< <i>System domain name</i> > | Randomly generated by the system | User for Loader system startup and Kerberos authentication                                                                             |                        |
|           | HTTP/< <i>hostname</i> >                    | Randomly generated by the system | Used to connect to the HTTP interface of each component. < <i>hostname</i> > indicates the host name of a node in the cluster.         |                        |
|           | hue                                         | Randomly generated by the system | User for Hue system startup, Kerberos authentication, and HDFS and Hive access                                                         |                        |
|           | flume                                       | Randomly generated by the system | User for Flume system startup and HDFS and Kafka access. The user has read and write permission of the HDFS directory / <b>flume</b> . |                        |
|           | flume_server                                | Randomly generated by the system | User for Flume system startup and HDFS and Kafka access. The user has read and write permission of the HDFS directory / <b>flume</b> . |                        |

| User Type | Default User                          | Initial Password                 | Description                                                                                                                                          | Password Change Method |
|-----------|---------------------------------------|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | spark2x/hadoop.<System domain name>   | Randomly generated by the system | This user is the Spark2x system administrator and has the following user permissions:<br>1. Starts the Spark2x service.<br>2. Submits Spark2x tasks. |                        |
|           | spark_zk/hadoop.<System domain name>  | Randomly generated by the system | Used for Spark2x to access ZooKeeper.                                                                                                                |                        |
|           | zookeeper/hadoop.<System domain name> | Randomly generated by the system | ZooKeeper system startup user.                                                                                                                       |                        |
|           | zkcli/hadoop.<System domain name>     | Randomly generated by the system | ZooKeeper server login user.                                                                                                                         |                        |
|           | oozie                                 | Randomly generated by the system | User for Oozie system startup and Kerberos authentication.                                                                                           |                        |
|           | kafka/hadoop.<System domain name>     | Randomly generated by the system | Used for security authentication of Kafka.                                                                                                           |                        |

| User Type | Default User                             | Initial Password                 | Description                                                                                                        | Password Change Method |
|-----------|------------------------------------------|----------------------------------|--------------------------------------------------------------------------------------------------------------------|------------------------|
|           | storm/hadoop.<System domain name>        | Randomly generated by the system | Storm system startup user.                                                                                         |                        |
|           | storm_zk/hadoop.<System domain name>     | Randomly generated by the system | Used for the Worker process to access ZooKeeper.                                                                   |                        |
|           | flink/hadoop.<System domain name>        | Randomly generated by the system | Internal user of the Flink service.                                                                                |                        |
|           | check_ker_M                              | Randomly generated by the system | User who performs a system internal test about whether the Kerberos service is normal.                             |                        |
|           | cdl/hadoop.<System domain name>          | Randomly generated by the system | Internal user of the CDL service.                                                                                  |                        |
|           | clickhouse / hadoop.<System domain name> | Randomly generated by the system | Used for security authentication of ClickHouse. This user is an internal user and can be used only in the cluster. |                        |
|           | default                                  | None                             | ClickHouse internal user, which is an administrator user that can be used only in non-security mode.               |                        |

| User Type | Default User                            | Initial Password                 | Description                                                                                                                                                                       | Password Change Method |
|-----------|-----------------------------------------|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|           | rangeradmin/hadoop.<System domain name> | Randomly generated by the system | Ranger system startup user, which is used for authentication between internal components.                                                                                         |                        |
|           | tez                                     | Randomly generated by the system | User for TezUI system startup, Kerberos authentication, and access to Yarn                                                                                                        |                        |
|           | K/M                                     | Randomly generated by the system | Kerberos internal functional user. It cannot be deleted, and its password cannot be changed. This internal account can only be used on nodes where Kerberos service is installed. | None                   |
|           | kadmin/changepw                         | Randomly generated by the system |                                                                                                                                                                                   |                        |
|           | kadmin/history                          | Randomly generated by the system |                                                                                                                                                                                   |                        |
|           | krbtgt<System domain name>              | Randomly generated by the system |                                                                                                                                                                                   |                        |
|           |                                         |                                  |                                                                                                                                                                                   |                        |



| User Type | Default User                          | Initial Password                     | Description                                                                                                                                                                                | Password Change Method                                                         |
|-----------|---------------------------------------|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| LDAP user | admin                                 | None                                 | FusionInsight Manager administrator.<br>The primary group is <b>compcommon</b> , which does not have the group permission but has the permission of the <b>Manager_administrator</b> role. | The LDAP user cannot log in to the system, and the password cannot be changed. |
|           | backup                                |                                      | The primary group is <b>compcommon</b> .                                                                                                                                                   |                                                                                |
|           | backup/manager                        |                                      | The primary group is <b>compcommon</b> .                                                                                                                                                   |                                                                                |
|           | oms                                   |                                      | The primary group is <b>compcommon</b> .                                                                                                                                                   |                                                                                |
|           | oms/manager                           |                                      | The primary group is <b>compcommon</b> .                                                                                                                                                   |                                                                                |
|           | clientregister                        |                                      | The primary group is <b>compcommon</b> .                                                                                                                                                   |                                                                                |
|           | zookeeper                             |                                      | The primary group is <b>hadoop</b> .                                                                                                                                                       |                                                                                |
|           | zookeeper/hadoop.<System domain name> |                                      | The primary group is <b>hadoop</b> .                                                                                                                                                       |                                                                                |
|           | zkcli                                 |                                      | The primary group is <b>hadoop</b> .                                                                                                                                                       |                                                                                |
|           | zkcli/hadoop.<System domain name>     |                                      | The primary group is <b>hadoop</b> .                                                                                                                                                       |                                                                                |
| flume     |                                       | The primary group is <b>hadoop</b> . |                                                                                                                                                                                            |                                                                                |

| User Type | Default User                       | Initial Password | Description                              | Password Change Method |
|-----------|------------------------------------|------------------|------------------------------------------|------------------------|
|           | flume_server                       |                  | The primary group is <b>hadoop</b> .     |                        |
|           | hdfs                               |                  | The primary group is <b>hadoop</b> .     |                        |
|           | hdfs/hadoop.<System domain name>   |                  | The primary group is <b>hadoop</b> .     |                        |
|           | mapred                             |                  | The primary group is <b>hadoop</b> .     |                        |
|           | mapred/hadoop.<System domain name> |                  | The primary group is <b>hadoop</b> .     |                        |
|           | mr_zk                              |                  | The primary group is <b>hadoop</b> .     |                        |
|           | mr_zk/hadoop.<System domain name>  |                  | The primary group is <b>hadoop</b> .     |                        |
|           | hue                                |                  | The primary group is <b>supergroup</b> . |                        |
|           | hive                               |                  | The primary group is <b>hive</b> .       |                        |
|           | hive/hadoop.<System domain name>   |                  | The primary group is <b>hive</b> .       |                        |
|           | hbase                              |                  | The primary group is <b>hadoop</b> .     |                        |
|           | hbase/hadoop.<System domain name>  |                  | The primary group is <b>hadoop</b> .     |                        |

| User Type | Default User                                                                    | Initial Password | Description                              | Password Change Method |
|-----------|---------------------------------------------------------------------------------|------------------|------------------------------------------|------------------------|
|           | thrift                                                                          |                  | The primary group is <b>hadoop</b> .     |                        |
|           | thrift/<br>hadoop.< <i>S</i><br><i>ystem</i><br><i>domain</i><br><i>name</i> >  |                  | The primary group is <b>hadoop</b> .     |                        |
|           | oozie                                                                           |                  | The primary group is <b>hadoop</b> .     |                        |
|           | hbase/<br>zkclient.< <i>S</i><br><i>ystem</i><br><i>domain</i><br><i>name</i> > |                  | The primary group is <b>hadoop</b> .     |                        |
|           | loader                                                                          |                  | The primary group is <b>hadoop</b> .     |                        |
|           | loader/<br>hadoop.< <i>S</i><br><i>ystem</i><br><i>domain</i><br><i>name</i> >  |                  | The primary group is <b>hadoop</b> .     |                        |
|           | spark2x                                                                         |                  | The primary group is <b>hadoop</b> .     |                        |
|           | spark2x/<br>hadoop.< <i>S</i><br><i>ystem</i><br><i>domain</i><br><i>name</i> > |                  | The primary group is <b>hadoop</b> .     |                        |
|           | spark_zk                                                                        |                  | The primary group is <b>hadoop</b> .     |                        |
|           | kafka                                                                           |                  | The primary group is <b>kafkaadmin</b> . |                        |
|           | kafka/<br>hadoop.< <i>S</i><br><i>ystem</i><br><i>domain</i><br><i>name</i> >   |                  | The primary group is <b>kafkaadmin</b> . |                        |
|           | storm                                                                           |                  | The primary group is <b>stormadmin</b> . |                        |

| User Type | Default User                            | Initial Password | Description                              | Password Change Method |
|-----------|-----------------------------------------|------------------|------------------------------------------|------------------------|
|           | storm/hadoop.<System domain name>       |                  | The primary group is <b>stormadmin</b> . |                        |
|           | storm_zk                                |                  | The primary group is <b>storm</b> .      |                        |
|           | storm_zk/hadoop.<System domain name>    |                  | The primary group is <b>storm</b> .      |                        |
|           | kms/hadoop                              |                  | The primary group is <b>kmsadmin</b> .   |                        |
|           | knox                                    |                  | The primary group is <b>compcommon</b> . |                        |
|           | executor                                |                  | The primary group is <b>compcommon</b> . |                        |
|           | rangeradmin                             |                  | The primary group is <b>supergroup</b> . |                        |
|           | rangeradmin/hadoop.<System domain name> |                  | The primary group is <b>supergroup</b> . |                        |
|           | rangerusersync                          |                  | The primary group is <b>supergroup</b> . |                        |
|           | rangertagsync                           |                  | The primary group is <b>supergroup</b> . |                        |
|           | rangerauditor                           |                  | The primary group is <b>compcommon</b> . |                        |
|           | jobserver                               |                  | The primary group is <b>compcommon</b> . |                        |

 **NOTE**

Log in to FusionInsight Manager, choose **System > Permission > Domain and Mutual Trust**, and check the value of **Local Domain**. In the preceding table, all letters in the system domain name contained in the username of the system internal user are lowercase letters.

For example, if **Local Domain** is set to **9427068F-6EFA-4833-B43E-60CB641E5B6C.COM**, the username of default HDFS startup user is **hdfs/hadoop.9427068f-6efa-4833-b43e-60cb641e5b6c.com**.

- **Database user**

The system database users include OMS database users and DBService database users.

| Database Type | Default User | Initial Password                                                                                                   | Description                                                                                                      | Password Change Method                                                                            |
|---------------|--------------|--------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| OMS database  | ommdba       | <ul style="list-style-type: none"> <li>• VerisiconsearlierthanMRSS3.2.0:dbchangeMe@123456</li> <li>• MR</li> </ul> | <p>OMS database administrator who performs maintenance operations, such as creating, starting, and stopping.</p> | <p>For details, see <a href="#">Changing the Password for the OMS Database Administrator</a>.</p> |

| Database Type | Default User | Initial Password                                                                                                                     | Description | Password Change Method |
|---------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------|
|               |              | S<br>3<br>.<br>2<br>.<br>0<br>o<br>r<br>l<br>a<br>t<br>e<br>r:<br>r<br>a<br>n<br>d<br>o<br>m<br>p<br>a<br>s<br>s<br>w<br>o<br>r<br>d |             |                        |

| Database Type | Default User | Initial Password                                                                                                                                                                                                                                                                                                                       | Description                          | Password Change Method                                                                  |
|---------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|-----------------------------------------------------------------------------------------|
|               | omm          | <ul style="list-style-type: none"> <li>• V<br/>e<br/>r<br/>s<br/>i<br/>o<br/>n<br/>s<br/>e<br/>a<br/>r<br/>l<br/>i<br/>e<br/>r<br/>t<br/>h<br/>a<br/>n<br/>M<br/>R<br/>S<br/>3<br/>.2<br/>.0<br/>:<br/>C<br/>h<br/>a<br/>n<br/>g<br/>e<br/>M<br/>e<br/>@<br/>1<br/>2<br/>3<br/>4<br/>5<br/>6</li> <li>• M<br/>R<br/>S<br/>3</li> </ul> | User for accessing OMS database data | For details, see <a href="#">Changing the Password for an OMS Database Access User.</a> |



| Database Type | Default User | Initial Password                                                                                                   | Description | Password Change Method |
|---------------|--------------|--------------------------------------------------------------------------------------------------------------------|-------------|------------------------|
|               |              | .2<br>.0<br>o<br>r<br>l<br>a<br>t<br>e<br>r:<br>r<br>a<br>n<br>d<br>o<br>m<br>p<br>a<br>s<br>s<br>w<br>o<br>r<br>d |             |                        |

| Database Type      | Default User | Initial Password                                                                                                                 | Description                                                      | Password Change Method                                                              |
|--------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| DBService database | omm          | <ul style="list-style-type: none"> <li>• Version earlier than MRSS3.2.0: <code>dbserverAdmin@123</code></li> <li>• MR</li> </ul> | Administrator of the GaussDB database in the DBService component | For details, see <a href="#">Resetting the Password for User omm in DBService</a> . |

| Database Type | Default User | Initial Password                                     | Description                                                                                                                                                                                    | Password Change Method                                                                                |
|---------------|--------------|------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
|               |              | S3.<br>2.<br>0<br>or<br>later:<br>random<br>password |                                                                                                                                                                                                |                                                                                                       |
|               | compdb user  | Random password                                      | (Available in MRS 3.1.2 or later) Administrator of the GaussDB database in the DBService component. It is used in service O&M scenarios. You need to reset the password upon your first login. | For details, see <a href="#">Changing the Password for User compdbuser of the DBService Database.</a> |

| Database Type | Default User | Initial Password | Description                                                                                                                | Password Change Method                                                                                                                                                                                                                                                                   |
|---------------|--------------|------------------|----------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | hetu         | Random password  | User for HetuEngine to connect to the DBService database <b>hetumeta</b> .<br>This user exists only in MRS 3.1.2 or later. | <ul style="list-style-type: none"> <li>For versions earlier than MRS 3.1.2, see <a href="#">Changing the Passwords for Database Users of MRS Cluster Components</a>.</li> <li>For MRS 3.1.2 or later, see <a href="#">Resetting the MRS Component Database User Password</a>.</li> </ul> |

| Database Type | Default User | Initial Password                                                                                                      | Description                                                          | Password Change Method |
|---------------|--------------|-----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|------------------------|
|               | hive         | <ul style="list-style-type: none"> <li>• Version earlier than MRSS3.1.2: HiveUser@</li> <li>• MRSS3.1.2 or</li> </ul> | User for Hive to connect to the DBService database <b>hivemeta</b> . |                        |

| Database Type | Default User | Initial Password                                                                             | Description | Password Change Method |
|---------------|--------------|----------------------------------------------------------------------------------------------|-------------|------------------------|
|               |              | l<br>a<br>t<br>e<br>r:<br>r<br>a<br>n<br>d<br>o<br>m<br>p<br>a<br>s<br>s<br>w<br>o<br>r<br>d |             |                        |

| Database Type | Default User | Initial Password                                                                                                      | Description                                            | Password Change Method |
|---------------|--------------|-----------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|------------------------|
|               | hue          | <ul style="list-style-type: none"> <li>• Version earlier than MRSS3.1.2: Hue User@123</li> <li>• MRSS3.1.2</li> </ul> | User for Hue to connect to the DBService database hue. |                        |

| Database Type | Default User | Initial Password                                                                                       | Description | Password Change Method |
|---------------|--------------|--------------------------------------------------------------------------------------------------------|-------------|------------------------|
|               |              | o<br>r<br>l<br>a<br>t<br>e<br>r:<br>r<br>a<br>n<br>d<br>o<br>m<br>p<br>a<br>s<br>s<br>w<br>o<br>r<br>d |             |                        |



| Database Type | Default User | Initial Password                                                                                                            | Description                                                         | Password Change Method |
|---------------|--------------|-----------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|------------------------|
|               | sqoop        | <ul style="list-style-type: none"> <li>• Version since earlier than MRSS3.1.2: SqoopUser@</li> <li>• MRSS3.1.2.0</li> </ul> | User for Loader to connect to the DBService database <b>sqoop</b> . |                        |

| Database Type | Default User | Initial Password                                                                                  | Description | Password Change Method |
|---------------|--------------|---------------------------------------------------------------------------------------------------|-------------|------------------------|
|               |              | r<br>l<br>a<br>t<br>e<br>r:<br>r<br>a<br>n<br>d<br>o<br>m<br>p<br>a<br>s<br>s<br>w<br>o<br>r<br>d |             |                        |

| Database Type | Default User | Initial Password                                                                                                              | Description                                                        | Password Change Method |
|---------------|--------------|-------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------|------------------------|
|               | oozie        | <ul style="list-style-type: none"> <li>• Version since earlier than MRSS3.1.2: Oozie User @</li> <li>• MRSS3.1.2.0</li> </ul> | User for Oozie to connect to the DBService database <b>oozie</b> . |                        |

| Database Type | Default User | Initial Password                                                                                  | Description | Password Change Method |
|---------------|--------------|---------------------------------------------------------------------------------------------------|-------------|------------------------|
|               |              | r<br>l<br>a<br>t<br>e<br>r:<br>r<br>a<br>n<br>d<br>o<br>m<br>p<br>a<br>s<br>s<br>w<br>o<br>r<br>d |             |                        |

| Database Type | Default User     | Initial Password                                                                                                                                                                                                                                                                                                                                 | Description                                           | Password Change Method |
|---------------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|------------------------|
|               | rangera<br>admin | <ul style="list-style-type: none"> <li>• V<br/>e<br/>r<br/>s<br/>i<br/>o<br/>n<br/>s<br/>e<br/>a<br/>r<br/>l<br/>i<br/>e<br/>r<br/>t<br/>h<br/>a<br/>n<br/>M<br/>R<br/>S<br/>3<br/>.<br/>1<br/>.<br/>2<br/>:<br/>A<br/>d<br/>m<br/>i<br/>n<br/>1<br/>2<br/>!</li> <li>• M<br/>R<br/>S<br/>3<br/>.<br/>1<br/>.<br/>2<br/>o<br/>r<br/>l</li> </ul> | User for Ranger to connect to the DBService database. |                        |

| Database Type | Default User   | Initial Password                                                                             | Description                                                                                                                              | Password Change Method |
|---------------|----------------|----------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
|               |                | a<br>t<br>t<br>e<br>r:<br>r<br>a<br>n<br>d<br>o<br>m<br>p<br>a<br>s<br>s<br>w<br>o<br>r<br>d |                                                                                                                                          |                        |
|               | kafkaui        | Ran<br>do<br>m<br>pass<br>wor<br>d                                                           | User for Kafka UI to connect to the DBService database.<br>This user exists only in MRS 3.1.2 or later.                                  |                        |
|               | flink          | Ran<br>do<br>m<br>pass<br>wor<br>d                                                           | User for Flink to connect to the DBService database <b>flinkmeta</b> .<br>This user exists only in MRS 3.1.2 or later.                   |                        |
|               | cdl            | Ran<br>do<br>m<br>pass<br>wor<br>d                                                           | User for CDL to connect to the DBService database <b>cdl</b> .<br>This user exists only in MRS 3.2.0 or later.                           |                        |
|               | jobgate<br>way | Ran<br>do<br>m<br>pass<br>wor<br>d                                                           | User for JobGateway to connect to the DBService database <b>jobmeta</b> .<br>This user can be used in MRS 3.3.0 and later versions only. |                        |

## Account List (MRS 2.x and Earlier Versions)

- **User type**

The MRS cluster provides the following three types of users. Periodically change the passwords and do not use the default passwords.

| User Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| System users          | <ul style="list-style-type: none"><li>• User created on Manager for MRS cluster O&amp;M and service scenarios. There are two types of system users:<ul style="list-style-type: none"><li>– <b>Human-machine</b> user: used for Manager O&amp;M scenarios and component client operation scenarios.</li><li>– <b>Machine-machine</b> user: used for MRS cluster application development scenarios.</li></ul></li><li>• User who runs OMS processes</li></ul> |
| Internal system users | Internal user to perform Kerberos authentication, process communications, save user group information, and associate user permissions. It is recommended that internal system users not be used in O&M scenarios. Operations can be performed as user <b>admin</b> or another user created by the system administrator based on service requirements.                                                                                                       |
| Database users        | <ul style="list-style-type: none"><li>• User who manages OMS database and accesses data</li><li>• User who runs the database of service components (Hive, Loader, and DBService)</li></ul>                                                                                                                                                                                                                                                                  |

- **System user**

 **NOTE**

- User **ldap** of the OS is required in the MRS cluster. Do not delete this account. Otherwise, the cluster may not work properly. Password management policies are maintained by the operation users.
- Reset the passwords when you change the passwords of user **ommdba** and user **omm** for the first time. Change the passwords periodically after retrieving them.

| User Type                               | Username | Initial Password                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------|----------|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| System administrator of the MRS cluster | admin    | Specified by the user during the cluster creation. | <p>MRS Manager</p> <p>The user has the following permissions:</p> <ul style="list-style-type: none"> <li>• Common HDFS and ZooKeeper user permissions.</li> <li>• Permissions to submit and query MapReduce and YARN tasks, manage YARN queues, and access the YARN web UI.</li> <li>• Permissions to submit, query, activate, deactivate, reassign, delete topologies, and operate all topologies of the Storm service.</li> <li>• Permissions to create, delete, authorize, reassign, consume, write, and query topics of the Kafka service.</li> </ul> |
| MRS cluster node OS user                | omm      | Randomly generated by the system                   | Internal running user of the MRS cluster system. This user is an OS user generated on all nodes and does not require a unified password.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| MRS cluster node OS user                | root     | The password is set by the user.                   | User for logging in to the node in the MRS cluster. This user is an OS user generated on all nodes.                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

- **Internal system user**

 **NOTE**

- Do not delete the following internal system users. Otherwise, the cluster or components may not work properly.
- Such user is available only in clusters with Kerberos authentication enabled.



| User Type              | Default User | Initial Password | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------|--------------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Component running user | hdfs         | Hdfs@123         | <p>This user is the HDFS system administrator and has the following permissions:</p> <ol style="list-style-type: none"><li>1. File system operation permissions:<ul style="list-style-type: none"><li>• Views, modifies, and creates files.</li><li>• Views and creates directories.</li><li>• Views and modifies the groups where files belong.</li><li>• Views and sets disk quotas for users.</li></ul></li><li>2. HDFS management operation permissions:<ul style="list-style-type: none"><li>• Views the web UI status.</li><li>• Views and sets the active and standby HDFS status.</li><li>• Enters and exits the HDFS in security mode.</li><li>• Checks the HDFS file system.</li></ul></li></ol> |

| User Type | Default User | Initial Password | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------|--------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | hbase        | Hbase@123        | <p>This user is the HBase system administrator and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Cluster management permission: <b>Enable</b> and <b>Disable</b> operations on tables to trigger MajorCompact and ACL operations.</li> <li>• Grants and revokes permissions, and shuts down the cluster.</li> <li>• Table management permission: Creates, modifies, and deletes tables.</li> <li>• Data management permission: Reads data in tables, column families, and columns.</li> <li>• Accesses the HBase web UI.</li> </ul> |
|           | mapred       | Mapred@123       | <p>This user is the MapReduce system administrator and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Submits, stops, and views the MapReduce tasks.</li> <li>• Modifies the Yarn configuration parameters.</li> <li>• Accesses the Yarn and MapReduce web UI.</li> </ul>                                                                                                                                                                                                                                                            |
|           | spark        | Spark@123        | <p>This user is the Spark system administrator and has the following permissions:</p> <ul style="list-style-type: none"> <li>• Accesses the Spark web UI.</li> <li>• Submits Spark tasks.</li> </ul>                                                                                                                                                                                                                                                                                                                                                              |

- **User group information**

| Default User Group    | Description                                                                                                                                                                                                                    |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| supergroup            | Primary group of user <b>admin</b> , which has no additional permissions in the cluster with Kerberos authentication disabled.                                                                                                 |
| check_sec_ldap        | Used to test whether the active LDAP works properly. This user group is generated randomly in a test and automatically deleted after the test is complete. This is an internal system user group used only between components. |
| Manager_tenant        | Tenant system user group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.                                                               |
| System_administrator  | MRS cluster system administrator group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.                                                 |
| Manager_viewer        | MRS Manager system viewer group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.                                                        |
| Manager_operator      | MRS Manager system operator group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.                                                      |
| Manager_auditor       | MRS Manager system auditor group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.                                                       |
| Manager_administrator | MRS Manager system administrator group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.                                                 |
| compcommon            | MRS cluster internal group, used to access public resources in the cluster. All system users and system running users are added to this user group by default.                                                                 |
| default_1000          | User group created for tenants. This is an internal system user group used only between components.                                                                                                                            |
| launcher-job          | MRS internal group, which is used to submit jobs using V2 APIs.                                                                                                                                                                |

| Default User Group | Description                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hadoop             | Users added to this user group have the permission to submit tasks to all YARN queues. Such user is available only in clusters with Kerberos authentication enabled.                                                                                            |
| hbase              | Common user group. Users added to this user group will not have any additional permission. Such user is available only in clusters with Kerberos authentication enabled.                                                                                        |
| hive               | Users added to this user group can use Hive. Such user is available only in clusters with Kerberos authentication enabled.                                                                                                                                      |
| spark              | Common user group. Users added to this user group will not have any additional permission. Such user is available only in clusters with Kerberos authentication enabled.                                                                                        |
| kafka              | Kafka common user group. Users added to this group need to be granted with read and write permission by users in the <b>kafkaadmin</b> group before accessing the desired topics. Such user is available only in clusters with Kerberos authentication enabled. |
| kafkasuperuser     | Users added to this group have permissions to read data from and write data to all topics. Such user is available only in clusters with Kerberos authentication enabled.                                                                                        |
| kafkaadmin         | Kafka administrator group. Users added to this group have the permissions to create, delete, authorize, as well as read from and write data to all topics. Such user is available only in clusters with Kerberos authentication enabled.                        |
| storm              | Storm common user group. Users added to this group have the permissions to submit topologies and manage their own topologies. Such user is available only in clusters with Kerberos authentication enabled.                                                     |
| stormadmin         | Storm administrator user group. Users added to this group have the permissions to submit topologies and manage their own topologies. Such user is available only in clusters with Kerberos authentication enabled.                                              |
| opentsdb           | Common user group. Users added to this user group will not have any additional permission. Such user is available only in clusters with Kerberos authentication enabled.                                                                                        |

| Default User Group | Description                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| presto             | Common user group. Users added to this user group will not have any additional permission. Such user is available only in clusters with Kerberos authentication enabled. |
| flume              | Common user group. Users added to this user group will not have any additional permission. Such user is available only in clusters with Kerberos authentication enabled. |
| launcher-job       | MRS internal group, which is used to submit jobs using V2 APIs. Such user is available only in clusters with Kerberos authentication enabled.                            |

| OS User Group | Description                                                                                                                          |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------|
| wheel         | Primary group of MRS internal running user <b>omm</b> .                                                                              |
| ficommon      | MRS cluster common group that corresponds to <b>comppcommon</b> for accessing public resource files stored in the OS of the cluster. |

- **Database user**

MRS cluster system database users include OMS database users and DBService database users.

 **NOTE**

Do not delete database users. Otherwise, the cluster or components may not work properly.

| User Type          | Default User | Initial Password  | Description                                                                                               |
|--------------------|--------------|-------------------|-----------------------------------------------------------------------------------------------------------|
| OMS database       | ommdba       | dbChangeMe@123456 | OMS database administrator who performs maintenance operations, such as creating, starting, and stopping. |
|                    | omm          | ChangeMe@123456   | User for accessing OMS database data                                                                      |
| DBService database | omm          | dbserverAdmin@123 | Administrator of the GaussDB database in the DBService component                                          |

| User Type | Default User | Initial Password | Description                                                                                                                         |
|-----------|--------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------|
|           | hive         | HiveUser@        | User for Hive to connect to the DBService database.                                                                                 |
|           | hue          | HueUser@123      | User for Hue to connect to the DBService database.                                                                                  |
|           | ranger       | RangerUser@      | User for Ranger to connect to the DBService database. Such user is available only in clusters with Kerberos authentication enabled. |
|           | sqoop        | SqoopUser@       | User for Loader to connect to the DBService database.                                                                               |

### 6.10.3 Managing MRS Cluster Roles

The administrator can create and manage different roles in Manager and use them to authorize access to Manager and components based on various service scenarios.

The Manager of MRS 3.x and later versions supports a maximum of 5,000 roles (including system built-in roles but excluding roles automatically created by tenants).

The Manager of MRS2.x and earlier versions supports a maximum of 1000 roles.

#### Creating a Role

**For MRS 3.x and later versions:**

**Step 1** Log in to FusionInsight Manager, choose **System > Permission > Role**.

**Step 2** On the displayed page, click **Create Role** and fill in **Role Name** and **Description**.

The role name consists of 3 to 50 characters, including digits, letters, and underscores (\_). It cannot be the same as an existing role name in the system. The role name cannot start with **Manager**, **System**, or **default**. For example, the role name cannot be **Manager\_test**.

**Figure 6-63** Creating a role

Role > Create Role

---

\* Role Name:

Configure Resource Permission:

| All resources |                    |
|---------------|--------------------|
| All resources | Description        |
| Manager       | Cluster Management |
| 312fts        |                    |

Description:

**Step 3** In the **Configure Resource Permission** area, click the cluster whose permissions are to be added and select service permissions for the role.

When setting permissions for a component, enter a resource name in the search text box in the upper right corner and click the search icon to view the search result.

The search result contains only directories, but not subdirectories. Search by keyword supports fuzzy match and is case-insensitive.

**NOTE**

- For components (except HDFS and Yarn) for which Ranger authorization has been enabled, the permissions of non-default roles on Manager do not take effect. You need to configure Ranger policies to assign permissions to user groups.
- If the resource requests of HDFS and Yarn are beyond the Ranger policies, the ACL rules of the components still take effect.
- A maximum of 1000 permissions can be set for a component at a time.

**Step 4** Click **OK**.

----End

**For MRS 2.x and earlier:**

**Step 1** On MRS Manager, choose **System > Manage Role**.


**Step 2** On the displayed page, click **Create Role** and fill in **Role Name** and **Description**.

**Role Name** is mandatory and contains 3 to 30 characters. Only digits, letters, and underscores (\_) are allowed. **Description** is optional.

**Step 3** In **Permission**, set role permission.

1. Click **Service Name** and select a name in **View Name**.
2. Select one or more permissions.

 **NOTE**

- The **Permission** parameter is optional.
- If you select **View Name** to set component permissions, you can enter a resource name in the **Search** box in the upper right corner and click . The search result is displayed.
- The search result contains only directories, but not subdirectories. Search by keyword supports fuzzy match and is case-insensitive. Results of the next page can be searched.

**Step 4** Click **OK**. Return to **Manage Role**.

----End

## Managing Roles

**Step 1** Log in to Manager.

**Step 2** Go to the role management page.

- For MRS 3.x and later versions, choose **System > Permission > Role**.
- For MRS 2.x and earlier versions, choose **System > Manage Role**.

**Step 3** In the role list, perform the following operations on the role as you need:

- To modify role information, locate the target role, and click **Modify** in the **Operation** column.
- To export role information, click **Export All** in **TXT** or **CSV** format. Role information includes the role name and description.

 **NOTE**

Exporting role information is available in MRS 3.x or later only.

- To delete a role, locate the row that contains the role and click **Delete**. To delete multiple roles in batches, select the target roles and click **Delete** above the role list. A role bound to a user cannot be deleted. To delete such a role, disassociate the role from the user by modifying the user first.

----End


## Example: Creating a Manager Role (MRS 3.x and Later)

**Step 1** Choose **System > Permission > Role**.

**Step 2** On the displayed page, click **Create Role** and fill in **Role Name** and **Description**.

**Step 3** In the **Configure Resource Permission** area, click **Manager** and set permissions for the role.

**Figure 6-64** Setting permissions



| Resource Name | Resource Type       | Permission                    |                                     |                                    |
|---------------|---------------------|-------------------------------|-------------------------------------|------------------------------------|
|               |                     | <input type="checkbox"/> view | <input type="checkbox"/> management | <input type="checkbox"/> Recursive |
| Cluster       | Manager access path | <input type="checkbox"/>      | <input type="checkbox"/>            |                                    |
| User          | Manager access path | <input type="checkbox"/>      | <input type="checkbox"/>            |                                    |
| Audit         | Manager access path |                               | <input type="checkbox"/>            |                                    |
| Tenant        | Manager access path |                               | <input type="checkbox"/>            |                                    |
| System        | Manager access path |                               | <input type="checkbox"/>            |                                    |



Manager permissions:

- Cluster
  - **view** permission: permission to view information on the **Cluster** page and view alarms and events under **O&M > Alarm**.
  - **management** permission: permission for management on the **Cluster** and **O&M** pages.
- User
  - **view** permission: permission to view information on pages under **System > Permission**.
  - **management** permission: permission for management on pages under **System > Permission**.
- Audit
  - management** permission: permission for management on the **Audit** page.
- Tenant
  - management** permission: permission for management on the **Tenant** page and permission to view alarms and events under **O&M > Alarm**.
- System
  - management** permission: permission for management on all pages except those under **Permission** on the **System** page and permission to view alarms and events under **O&M > Alarm**.

**Step 4** Click **OK**.

----End

## 6.10.4 Managing MRS Cluster User Groups

Administrators can create and manage different user groups based on service scenarios on Manager. A user group is bound to a role to obtain operation permissions. After a user is added to a user group, the user can obtain the operation permissions of the user group. A user group can be used to classify users and manage multiple users.

The Manager of an MRS 3.x or later cluster supports a maximum of 5000 user groups (including built-in user groups).

The Manager of an MRS 2.x or earlier cluster supports a maximum of 100 user groups (including built-in user groups).

### Prerequisites

- You have learned service requirements and created roles required by service scenarios.
- You have logged in to Manager.

### Creating a User Group

**For MRS 3.x and later versions:**

**Step 1** Choose **System > Permission > User Group**.

**Step 2** Above the user group list, click **Create User Group**.

**Figure 6-65** Creating a user group

User Group > **Create User Group**

---

\* Group Name:

Role: [Add](#) [Clear All](#)

User: [Add](#) [Clear All](#)

Description:

**Step 3** Set **Group Name** and **Description**.

The group name contains 1 to 64 characters, including case-insensitive letters, digits, underscores (\_), hyphens (-), and spaces. It cannot be the same as an existing user group name in the system.

**Step 4** In the **Role** area, click **Add** to select a role and add it.

 **NOTE**

- For components (except HDFS and Yarn) for which Ranger authorization has been enabled, the permissions of non-default roles on Manager do not take effect. You need to configure Ranger policies to assign permissions to user groups.
- If the resource requests of HDFS and Yarn are beyond the Ranger policies, the ACL rules of the components still take effect.

**Step 5** In the **User** area, click **Add** to select a user and add it.

**Step 6** Click **OK**.

The user group is created.

----**End**

**For MRS 2.x and earlier:**

**Step 1** On MRS Manager, click **System**.

**Step 2** In the **Permission** area, click **Manage User Group**.

**Step 3** Above the user group list, click **Create User Group**.

**Step 4** Set **Group Name** and **Description**.

**Group Name** is mandatory and contains 3 to 20 characters. Only digits, letters, and underscores (\_) are allowed. **Description** is optional.

**Step 5** In **Role**, click **Select and Add Role** to select and add specified roles.

If you do not add the roles, the user group you are creating now does not have the permission to use MRS clusters.

**Step 6** Click **OK**.

----End

## Managing User Groups

**Step 1** Log in to Manager.

**Step 2** Go to the user group management page.

- For MRS 3.x and later versions, choose **System > Permission > User Group**.
- For MRS 2.x and earlier versions, Click **System**. In the **Permission** area, click **Manage User Group**.

**Step 3** In the user group list, perform the following operations on a user group as you need.

- Viewing user group information: By default, all user groups are displayed in the user group list. You can click the arrow on the left of a user group name to view details about the user group, including the user quantity, specific users, and bound roles of the user group.
- Modifying user group information: Locate the row that contains the user group to be modified and click **Modify**.
- Exporting user group information: Click **Export All** to export all user group information at a time in **TXT** or **CSV** format.

The user group information contains the user group name, description, user list, and role list.

### NOTE

Exporting user group information is supported in MRS 3.x or later only.

- Deleting a user group: Locate the row that contains the target user group and click **Delete**. To delete multiple user groups in batches, select the target user groups and click **Delete** above the user group list. A user group that contains users cannot be deleted. To delete such a user group, delete all its users by modifying the user group first.

----End

## 6.10.5 Managing MRS Cluster Users

### 6.10.5.1 Creating an MRS Cluster User

By default, only user **admin** has the highest operation permissions of Manager. Administrators need to create users on Manager and assign operation permissions to the users based on service requirements.

The FusionInsight Manager of an MRS 3.x or later cluster supports a maximum of 50000 users (including built-in users).

The MRS Manager of an MRS 2.x or earlier cluster supports a maximum of 1000 users.

#### Creating a User (MRS 3.x and Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **System > Permission > User**.

**Step 3** On the **User** page, click **Create**.

**Step 4** Set **Username**. The username can contain digits, letters, underscores (`_`), hyphens (`-`), and spaces. It is case-insensitive and cannot be the same as any existing username in the system or OS.

**Step 5** Set **User Type** to **Human-Machine** or **Machine-Machine**.

- **Human-Machine** user: used for FusionInsight Manager O&M and component client operations. If you select this option, you also need to select the password policy and set **Password** and **Confirm Password**.
- **Machine-Machine** user: used for component application development. If you select this option, the password is randomly generated.

**Step 6** In the **User Group** area, click **Add** to add one or more user groups to the list.

#### NOTE

- If the selected user group has been bound to a role or a permission policy has been configured in Ranger, the user can obtain the corresponding permissions.
- After FusionInsight Manager is installed, some user groups generated by default have special permissions. Select desired user groups based on the descriptions on the UI.
- If existing user groups cannot meet your requirements, click **Create User Group** to create a user group. For details, see [Creating a User Group](#).

**Step 7** Select a group from the **Primary Group** drop-down list to create directories and files.

The drop-down list contains all groups selected in **User Group**.

#### NOTE

A user can belong to multiple groups (including the primary group and secondary groups). The primary group is set to facilitate maintenance and comply with the permission mechanism of the Hadoop community. The primary group has the same permission control functionality as other groups.

**Step 8** In the **Role** area, click **Add** to bind roles to the user.

 NOTE

- Adding a role when you create a user can specify the user permissions.
- If the permissions granted to the user from the user group cannot meet service requirements, you can bind other created roles to the user. You can also click **Create Role** to create a role first. For details, see [Creating a Role](#).  
It takes 3 minutes to make role permission assignment to the user take effect. If the permissions obtained from the user group are enough, you do not need to add a role.
- After Ranger authentication is enabled for a component, you need to configure Ranger policies to assign permissions to the user except the permissions of default user group or role.
- If a user is not added to a user group or assigned a role, the user cannot view information or perform operations after logging in to FusionInsight Manager.

**Step 9** Enter information in **Description**.

**Step 10** Click **OK**.

After a human-machine user is created, you need to change the initial password as prompted after logging in to FusionInsight Manager.

----End

## Creating a User (MRS 2.x and Earlier)

**Step 1** On MRS Manager, click **System**.

**Step 2** In the **Permission** area, click **Manage User**.

**Step 3** On the **User** page, click **Create**.

**Step 4** Configure parameters as prompted and enter a username in **Username**.

 NOTE

- A username that differs only in alphabetic case from an existing username is not allowed. For example, if **User1** has been created, you cannot create **user1**.
- When you use the user you created, enter the exactly correct username, which is case-sensitive.
- **Username** is mandatory and contains 3 to 20 characters. Only digits, letters, and underscores (\_) are allowed.
- **root**, **omm**, and **ommdba** are reserved system user. Select another username.

**Step 5** Set **User Type** to **Human-Machine** or **Machine-Machine**.

- **Human-machine** user: used for MRS Manager O&M scenarios and component client operation scenarios. If you select this user type, you need to enter a password and confirm the password in **Password** and **Confirm Password** accordingly.
- **Machine-machine** users: used for MRS application development scenarios. If you select this user type, you do not need to enter a password, because the password is randomly generated.

**Step 6** In **User Group**, click **Select and Join User Group** to select user groups and add users to them.

 **NOTE**

- If roles have been added to user groups, the users can be granted with permissions of the roles.
- If you want to grant new users with Hive permissions, add the users to the Hive group.
- If a user needs to manage tenant resources, the user group must be assigned the **Manager\_tenant** role and the role corresponding to the tenant.

**Step 7** In **Primary Group**, select a group as the primary group for users to create directories and files. The drop-down list contains all groups selected in **User Group**.

**Step 8** In **Assign Rights by Role**, click **Select and Add Role** to add roles for users based on onsite service requirements.

 **NOTE**

- If the permissions granted to the user from the user group cannot meet service requirements, you can bind other created roles to the user. It takes 3 minutes to make role permissions granted to the new user take effect.
- Adding a role when you create a user can specify the user permissions.
- A new user can access web UIs of HDFS, HBase, Yarn, Spark, and Hue even when roles are not assigned to the user.

**Step 9** In **Description**, provide description based on onsite service requirements.

**Description** is optional.

**Step 10** Click **OK**.

If a new user is used in the MRS cluster for the first time, for example, used for logging in to MRS Manager or using the cluster client, the password must be changed.

----End

### 6.10.5.2 Modifying MRS Cluster User Information

You can modify user information on Manager, including the user group, primary group, role permission assignment, and user description.

#### Modifying User Information (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **System > Permission > User**.

**Step 3** Locate the row that contains the target user and click **Modify** in the **Operation** column.

Modify the parameters based on service requirements.

 **NOTE**

It takes three minutes at most for the change of the user group or role permissions to take effect.

MRS 3.1.2 or later:

- Users (except **admin**) cannot modify their own password policies.
- Locked users cannot modify their password policies.
- After the password policy bound to a user is modified, the modification takes effect when the user changes the password next time.
- After the password policy bound to a user is modified, if the remaining password validity period is greater than the password validity period in the new password policy, the password validity period is set to the validity period in the new password policy. If the remaining password validity period is less than the password validity period in the new password policy, the password validity period remains unchanged.

**Step 4** Click **OK**.

----End

## Modifying User Information (MRS 2.x and Earlier)

**Step 1** On MRS Manager, click **System**.

**Step 2** In the **Permission** area, click **Manage User**.

**Step 3** In the row of a user to be modified, click **Modify**.

 **NOTE**

If you change user groups for a user or assign role permissions to a user, it takes 3 minutes to apply the new configurations.

**Step 4** Click **OK**. The modification is complete.

----End

### 6.10.5.3 Locking an MRS Cluster User

A user may be suspended for a long period of time due to service changes. For security purposes, you can lock such a user. You can lock a user in using either of the following methods:

- Automatic locking: You can set **Password Retries** in the password policy to automatically lock the user whose login attempts exceed this parameter value. For details, see [Configuring Password Policies for MRS Cluster Users](#).
- Manual locking: You manually lock a user.

This section describes how to lock a user manually. Machine-machine users cannot be locked. For details about how to unlock an MRS cluster user, see [Unlocking a User Created on Manager](#).

## Impact on the System

A locked user cannot log in to Manager or perform identity authentication in the cluster. A locked user can be used only after being manually unlocked or the lock time expires.

## Locking a User

**For MRS 3.x and later versions:**

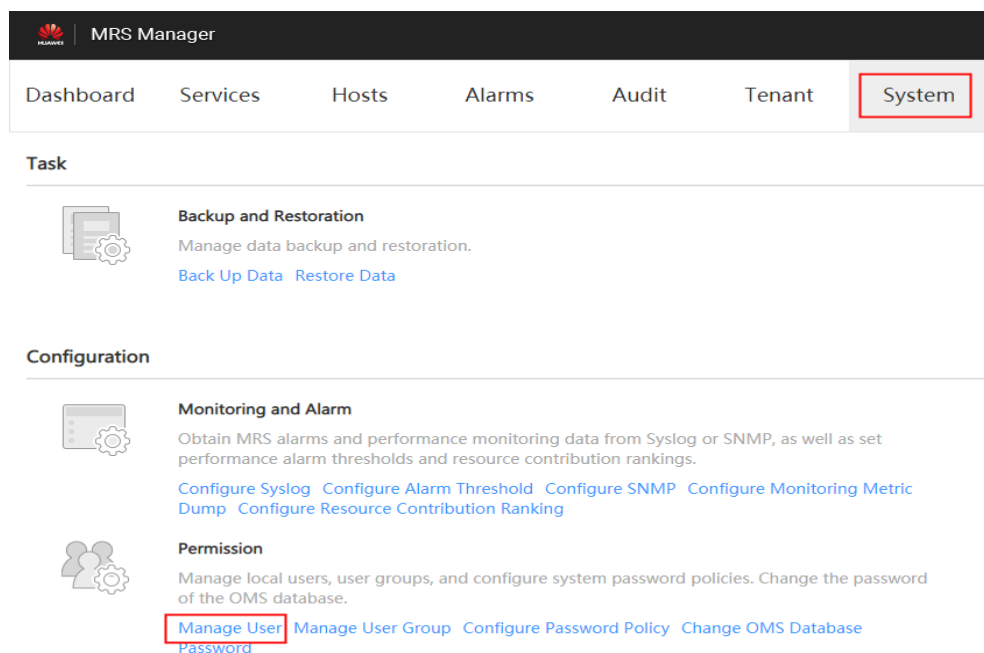
- Step 1** Log in to Manager.
- Step 2** Choose **System > Permission > User**.
- Step 3** Locate the row that contains the target user and click **Lock** in the **Operation** column.
- Step 4** In the window that is displayed, select **I have read the information and understand the impact**. Click **OK**.

----End

**For MRS 2.x and earlier:**

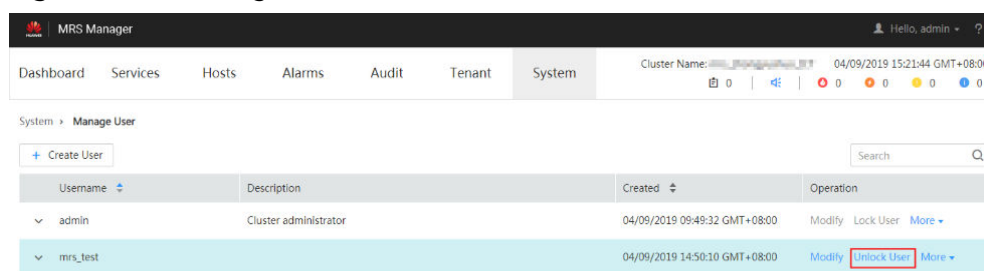
- Step 1** On MRS Manager, click **System**.
- Step 2** In the **Permission** area, click **Manage User**.

**Figure 6-66** User management



- Step 3** In the row of a user you want to lock, click **Lock User**.

**Figure 6-67** Locking a user





**Step 4** In the window that is displayed, click **OK** to lock the user.

----End

### 6.10.5.4 Deleting an MRS Cluster User

Based on service requirements, you can delete system users that are no longer used on Manager.

 **NOTE**

- After a user is deleted, the provisioned ticket granting ticket (TGT) is still valid within 24 hours. The user can use the TGT for security authentication and access the system.
- If a new user has the same name as the deleted user, the new user will inherit all owner permissions of the deleted user. You are advised to determine whether to delete the resources owned by the deleted user based on service requirements, for example, files in HDFS.
- The default user **admin** cannot be deleted.

### Deleting a Cluster User (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **System > Permission > User**.

**Step 3** Locate the row that contains the target user, click **More**, and select **Delete**.

 **NOTE**

To delete users in batches, select the users at a time and click **Delete**.

**Step 4** In the displayed dialog box, click **OK**.

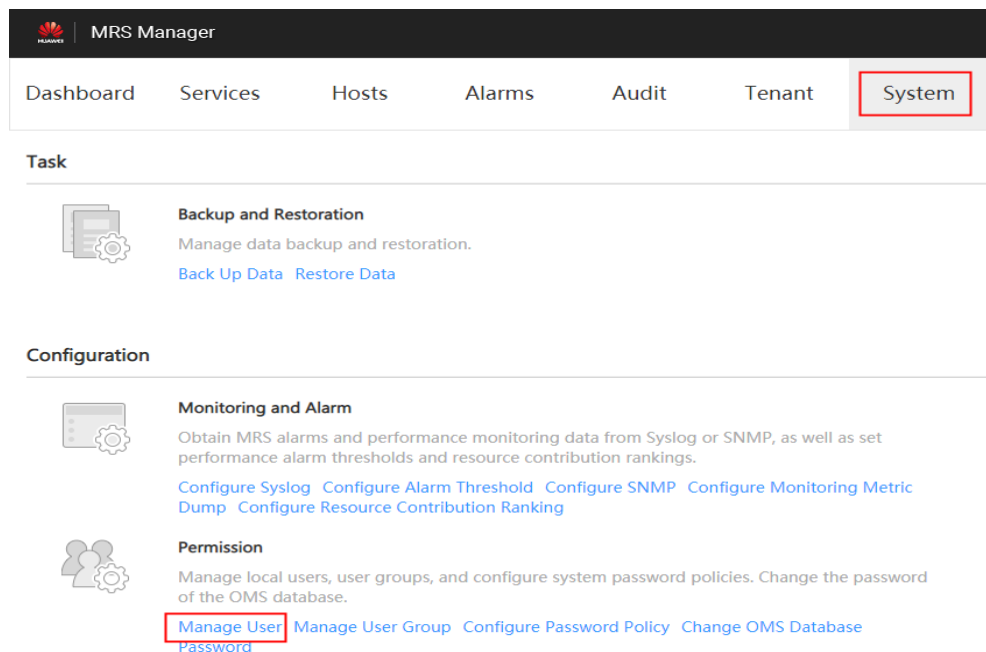
----End

### Deleting a Cluster User (MRS 2.x and Earlier)

**Step 1** On MRS Manager, click **System**.


**Step 2** In the **Permission** area, click **Manage User**.

**Figure 6-68** User management



**Step 3** In the row that contains the user to be deleted, choose **More > Delete**.

**Figure 6-69** Deleting a user



**Step 4** Click **OK**.

**NOTE**

If you want to create a new user with the same name as user A after deleting user A who has submitted a job on the client or MRS console, you need to delete user A's residual folders when deleting user A. Otherwise, the newly created user A may fail to submit a job.

To delete residual folders, log in to each Core node in the MRS cluster and run the following commands. In the following commands, **\$user** indicates the folder named after the username.

```
cd /srv/BigData/hadoop/data1/nm/localdir/usercache/
rm -rf $user
```

----End

### 6.10.5.5 Initializing MRS Cluster User Passwords

To reset a forgotten password or periodically change a public account password, administrator can initialize a new password on the Manager. The system user must then change the password upon their first login.

#### Impact on the System

If you have downloaded a user authentication file, download it again and obtain the keytab file after initializing the password of the MRS cluster user.

#### Initializing the Password of a Human-Machine User

**For MRS 3.x and later versions:**

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **System > Permission > User**.
- Step 3** Locate the row that contains the target user, click **More**, and select **Initialize Password**. In the displayed dialog box, enter the password of the current login user and click **OK**. In the **Initialize Password** dialog box, click **OK**.
- Step 4** Set **New Password** and **Confirm Password**, and click **OK**.

The password must meet the following complexity requirements:

- Contains at least 8 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (~!@#%&\*()-\_+=|[{]}';<.>/\?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be a common easily-cracked password.
- Cannot be the same as the password used in the latest *N* times. *N* indicates the value of **Repetition Rule** configured in [Configuring Password Policies for MRS Cluster Users](#).

----End

**For MRS 2.x and earlier:**

- Step 1** On MRS Manager, click **System**.
- Step 2** In the **Permission** area, click **Manage User**.
- Step 3** Locate the row that contains the user whose password is to be initialized, choose **More > Initialize password**, and change the password as prompted.

In the window that is displayed, enter the password of the current administrator account and click **OK**. Then in **Initialize password**, click **OK**.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters (~!@#%&\*()-\_+=\|[{]}';<.>/\?).

- The password cannot be the username or the reverse username.

----End

## Initializing the Password of a Machine-Machine User

**Step 1** Prepare a client based on service conditions and log in to the node with the client installed.

**Step 2** Run the following command to switch the user:

```
sudo su - omm
```

**Step 3** Run the following command to switch to the client directory, for example, **/opt/client**:

```
cd /opt/client
```

**Step 4** Run the following command to set environment variables:

```
source bigdata_env
```

**Step 5** Run the following command to log in to the console as user **kadmin/admin**:

```
kadmin -p kadmin/admin
```

### NOTE

The default password of user **kadmin/admin** is **KAdmin@123**, which will expire upon your first login. Change the password as prompted and keep the new password secure.

**Step 6** Run the following command to reset the password of a component running user. This operation takes effect on all servers:

```
cpw Component running user name
```

For example, **cpw oms/manager**.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#\$%^&\*()-\_+=\|[]{};:'",.<.>/?).
- The password cannot be the username or the reverse username.

----End

### 6.10.5.6 Downloading MRS Cluster User Credentials

If you develop big data applications and run them in an MRS cluster that requires Kerberos authentication, you need to prepare a user authentication file for accessing the MRS cluster. The keytab file in the authentication files is used for user authentication.

This topic describes how to download user authentication files and export the keytab file on MRS Manager.

 NOTE

After a user password is changed, the exported keytab file becomes invalid, and you need to export a keytab file again.

## Prerequisites

Before downloading the keytab file of a Human-Machine user, the password of the user must be changed at least once on the Manager portal or a client; otherwise, the downloaded keytab file cannot be used. For details, see [Changing the Passwords for Manager Users of an MRS Cluster](#).

## Downloading the Authentication Credential

**Step 1** Log in to Manager.

For MRS 3.x and later versions, choose **System > Permission > User**.

For versions earlier than MRS 3.x, choose **System > Permission > Manage User**.

**Step 2** Locate the row that contains the user whose keytab file needs to be exported, choose **More > Download Authentication Credential**, specify the save path after the file is automatically generated, and keep the file properly.

The authentication credential includes the **krb5.conf** file of the Kerberos service.

After the authentication credential file is decompressed, you can obtain the following two files:

- The **krb5.conf** file contains the authentication service connection information.
- The **user.keytab** file contains user authentication information.

----End

## 6.10.6 Unlocking an MRS Cluster User

### 6.10.6.1 Unlocking an LDAP User in the MRS Cluster

To unlock a user who has been locked out due to excessive incorrect password attempts, log in to the MRS cluster node and run commands to unlock the user, or use Manager to unlock the user if they were created on Manager.

If the service is abnormal, the internal user of the system may be locked. Unlock the user promptly, or the cluster cannot run properly. For the internal user list, see [MRS Cluster User Accounts](#). Internal users cannot be unlocked on Manager.

## Unlocking a User Created on Manager

**For MRS 3.x and later versions:**

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **System > Permission > User**.

**Step 3** Locate the row that contains the target user and click **Unlock** in the **Operation** column.

**Step 4** In the window that is displayed, select **I have read the information and understand the impact**. Click **OK**.


----End

**For MRS 2.x and earlier:**

**Step 1** On MRS Manager, click **System**.


**Step 2** In the **Permission** area, click **Manage User**.

**Figure 6-70** User management



**Step 3** In the row of a user to be unlocked, click **Unlock User**.

**Figure 6-71** Unlocking a user



**Step 4** In the window that is displayed, click **OK** to unlock the user.

----End

## Unlocking an Internal User

This function is available in MRS 3.x and later only.

**Step 1** Use the following method to confirm whether the internal system username is locked:

1. OLdap port number obtaining method:
  - a. Log in to Manager, choose **System > OMS > oldap > Modify Configuration**.
  - b. The **LDAP Listening Port** parameter value is **oldap port**.
2. Domain name obtaining method:
  - a. Log in to Manager, choose **System > Permission > Domain and Mutual Trust**.
  - b. The **Local Domain** parameter value is the domain name.  
For example, the domain name of the current system is **9427068F-6EFA-4833-B43E-60CB641E5B6C.COM**.

3. Run the following command on each node in the cluster as user **omm** to query the number of password authentication failures:

```
ldapsearch -H ldaps://OMS Floating IP Address:Oldap port -LLL -x -D
cn=root,dc=hadoop,dc=com -b krbPrincipalName=Internal system
username@Domain name,cn=Domain
name,cn=krbcontainer,dc=hadoop,dc=com -w Password of LDAP
administrator -e ppolicy | grep krbLoginFailedCount
```

#### NOTE

- To obtain the floating IP address of OMS, log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of OMS. Record the value of **inet**. If the floating IP address of OMS cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one Master node, query and record the cluster manager IP address of the Master node.
- LDAP administrator password: Obtain the default password of LDAP administrator **cn=root,dc=hadoop,dc=com** by referring to [MRS Cluster User Accounts](#).

For example, run the following command to check the number of password authentication failures for user **oms/manager**:

```
ldapsearch -H ldaps://10.5.146.118:21750 -LLL -x -D
cn=root,dc=hadoop,dc=com -b krbPrincipalName=oms/
manager@9427068F-6EFA-4833-
B43E-60CB641E5B6C.COM,cn=9427068F-6EFA-4833-
B43E-60CB641E5B6C.COM,cn=krbcontainer,dc=hadoop,dc=com -w
Password of user cn=root,dc=hadoop,dc=com -e ppolicy | grep
krbLoginFailedCount
```

```
krbLoginFailedCount: 5
```

4. Log in to Manager, choose **System > Permission > Security Policy > Password Policy**.
5. Check the value of the **Password Retries** parameter. If the value is less than or equal to the value of **krbLoginFailedCount**, the user is locked.

#### NOTE

You can also check whether internal users are locked by viewing operations logs.

- Step 2** Log in to the active management node as user **omm** and run the following command to unlock the user:

```
sh ${BIGDATA_HOME}/om-server/om/share/om/acs/config/unlockuser.sh --
userName Internal system username
```

```
Example: sh ${BIGDATA_HOME}/om-server/om/share/om/acs/config/
unlockuser.sh --userName oms/manager
```

----End

### 6.10.6.2 Unlocking the LDAP Management Account of the MRS Cluster

If the LDAP user **cn=pg\_search\_dn,ou=Users,dc=hadoop,dc=com** and LDAP management accounts **cn=krbkdc,ou=Users,dc=hadoop,dc=com** and **cn=krbadmin,ou=Users,dc=hadoop,dc=com** are locked, the administrator must unlock these accounts.

#### NOTE

- If you input an incorrect password for the LDAP user or management account for five consecutive times, the LDAP user or management account is locked. The account is automatically unlocked after 5 minutes.
- This function is available in MRS 3.x and later only.

**Step 1** Log in to the active management node as user **omm**.

**Step 2** Run the following command to go to the related directory:

```
cd ${BIGDATA_HOME}/om-server/om/ldapserver/ldapserver/local/script
```

**Step 3** Run the following command to unlock the LDAP user or management account:

```
./ldapserver_unlockUsers.sh USER_NAME
```

In the command, *USER\_NAME* indicates the name of the user to be unlocked.

For example, to unlock the LDAP management **account** **cn=krbkdc,ou=Users,dc=hadoop,dc=com**, run the following command:

```
./ldapserver_unlockUsers.sh krbkdc
```

After the script is executed, enter the password of user **krbkdc** next to **ROOT\_DN\_PASSWORD**. If the following information is displayed, the unlocking is successful:

```
Unlock user krbkdc successfully.
```

----End

### 6.10.7 Configuring Password Policies for MRS Cluster Users

To keep up with service security requirements, you can set password security rules, user login security rules, and user locking rules on Manager.

#### NOTE

- Modify password policies based on service security requirements, because they involve user management security. Otherwise, security risks may be incurred.
- Change the user password after modifying the password policy, and then the new password policy can take effect.
- This password policy is used for human-machine users created on MRS. A maximum of 32 password policies can be created.



## Adding a Password Policy

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **System > Permission > Security Policy > Password Policy**.

**Step 3** Click **Add Password Policy** and modify the password policy as prompted.

For details about the parameters, see [Table 6-59](#).

**Table 6-59** Password policy parameters

| Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password Policy Name           | The value is a string of 3 to 32 characters, including case-insensitive letters, digits, underscores (_), and hyphens (-). It cannot start with a hyphen (-).                                                                                                                                                                                                                                                                                                                                                 |
| Minimum Password Length        | Indicates the minimum number of characters a password contains. The value ranges from <b>8</b> to <b>32</b> .                                                                                                                                                                                                                                                                                                                                                                                                 |
| Character Types                | Indicates how many character types in the following types a password can contain at least: uppercase letters, lowercase letters, digits, spaces, and special characters (~!?,,;:_'(){}[]/<>@#\$\$%^&*+ \=). The value can be <b>4</b> or <b>5</b> . The default value is <b>4</b> , which means that a password can contain uppercase letters, lowercase letters, digits, and special characters. If you set the parameter to <b>5</b> , a password can contain all the five character types mentioned above. |
| Password Retries               | Indicates the number of consecutive wrong password attempts allowed before the system locks the user. The value ranges from <b>3</b> to <b>30</b> .                                                                                                                                                                                                                                                                                                                                                           |
| User Lock Duration (Min)       | Indicates the time period in which a user is locked when the user lockout conditions are met. The value ranges from <b>5</b> to <b>120</b> .                                                                                                                                                                                                                                                                                                                                                                  |
| Password Validity Period (Day) | Indicates the validity period of a password. The value ranges from <b>0</b> to <b>90</b> . <b>0</b> indicates that the password is permanently valid.                                                                                                                                                                                                                                                                                                                                                         |
| Repetition Rule                | Indicates the number of previous passwords that cannot be reused when you change the password. The value ranges from <b>1</b> to <b>5</b> . The default value is <b>1</b> .<br>This policy applies to only human-machine accounts.                                                                                                                                                                                                                                                                            |

| Parameter                                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password Expiration Notification (Days)                    | Indicates the number of days in advance users are notified that their passwords are about to expire. After the value is set, if the difference between the cluster time and the password expiration time is smaller than this value, the user receives password expiration notifications. When a user logs in to Manager, a message is displayed, indicating that the password is about to expire and asking the user whether to change the password. The value ranges from <b>0</b> to <i>X</i> ( <i>X</i> must be set to the half of the password validity period and rounded down). Value <b>0</b> indicates that no notification is sent. The default value is <b>5</b> . |
| Interval for Deleting Authentication Failure Records (Min) | Indicates the interval of retaining incorrect password attempts. The value ranges from <b>0</b> to <b>1440</b> . <b>0</b> indicates that incorrect password attempts are permanently retained, and <b>1440</b> indicates that incorrect password attempts are retained for one day.                                                                                                                                                                                                                                                                                                                                                                                           |

**Step 4** Click **OK** to save the configurations.

A new user uses the default password policy. After a new password policy is created, you can manually select the password policy when creating a user. You can modify the password policy of an existing user. For details, see [Modifying MRS Cluster User Information](#).

**Step 5** To delete a manually added password policy, perform the following operations:

Click **Delete** in the row that contains the target password policy. In the dialog box that is displayed, click **OK**.

 **NOTE**

The default password policy and the password policy that has been bound to a user cannot be deleted.

----End

## Modifying a Password Policy

**Step 1** Log in to Manager.

**Step 2** Enter the password policy configuration page.

- For MRS 2.x and earlier versions: Choose **System > Password Policy Configuration**.
- For MRS 3.x and later versions, choose **System > Permission > Security Policy > Password Policy**, and click **Modify** in the row that contains the password policy you want to modify.

**Step 3** Modify password policies as prompted.

**Table 6-60** Password policy parameters

| Parameter                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Minimum Password Length                 | Indicates the minimum number of characters a password contains. The value ranges from 8 to 32. The default value is <b>8</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Character Types                         | Indicates how many character types in the following types a password can contain at least: uppercase letters, lowercase letters, digits, spaces, and special characters (~`!?,,:;-'(){}[]/<>@#\$\$%^&*+ \=). <ul style="list-style-type: none"> <li>For MRS 3.x and later versions, the value can be <b>4</b> or <b>5</b>. The default value is <b>4</b>, which means that a password can contain uppercase letters, lowercase letters, digits, and special characters. If you set the parameter to <b>5</b>, a password can contain all the five character types mentioned above.</li> <li>For MRS2.x and earlier versions, the value can be <b>3</b> or <b>4</b>. The default value <b>3</b> indicates that the password must contain at least three types of the following characters: uppercase letters, lowercase letters, digits, special characters, and spaces.</li> </ul> |
| Password Retries                        | Indicates the number of consecutive wrong password attempts allowed before the system locks the user. The value ranges from <b>3</b> to <b>30</b> . The default value is <b>5</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| User Lock Duration (Min)                | Indicates the time period in which a user is locked when the user lockout conditions are met. The value ranges from <b>5</b> to <b>120</b> . The default value is <b>5</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Password Validity Period (Day)          | Indicates the validity period (days) of a password. The value ranges from 0 to 90. Value <b>0</b> means that the password is permanently valid. The default value is <b>90</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Repetition Rule                         | Indicates the number of previous passwords that cannot be reused when you change the password. The value ranges from <b>1</b> to <b>5</b> . The default value is <b>1</b> . This parameter is required for clusters of MRS 3.x or later. This policy applies to only human-machine accounts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Password Expiration Notification (Days) | Indicates the number of days in advance users are notified that their passwords are about to expire. After the value is set, if the difference between the cluster time and the password expiration time is smaller than this value, the user receives password expiration notifications. When a user logs in to Manager, a message is displayed, indicating that the password is about to expire and asking the user whether to change the password. The value ranges from <b>0</b> to <i>X</i> ( <i>X</i> must be set to the half of the password validity period and rounded down). Value <b>0</b> indicates that no notification is sent. The default value is <b>5</b> .                                                                                                                                                                                                      |

| Parameter                                                  | Description                                                                                                                                                                                                                                                                                                    |
|------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interval for Deleting Authentication Failure Records (Min) | Indicates the interval (minutes) of retaining incorrect password attempts. The value ranges from 0 to 1440. <b>0</b> indicates that incorrect password attempts are permanently retained, and <b>1440</b> indicates that incorrect password attempts are retained for one day. The default value is <b>5</b> . |

**Step 4** Click **OK** to save the configurations. Change the user password after modifying the password policy, and then the new password policy can take effect.

 **NOTE**

For MRS 3.1.2 and later:

- Users (except **admin**) cannot modify their own password policies.
- When a user's password policy is updated, the password's remaining validity period is adjusted as follows: if it is longer than the new policy's validity period, it is shortened to match the new policy; otherwise, it remains unchanged.

----End

## 6.10.8 Configuring the Private Attribute of MRS Cluster Users

User **admin** or administrators who are bound to the Manager\_administrator role can configure the independent attribute on Manager so that common users (all service users in the cluster) can set or cancel their own independent attributes.

After the independent attribute option is toggled on, service users need to log in to the system and set the independent attribute.

 **NOTE**

This topic is available for MRS 3.x and later only.

### Constraints

- Administrators cannot set or cancel the independent attribute of a user.
- Administrators cannot obtain the authentication credentials of independent users.

### Prerequisites

You have obtained the required administrator username and password.

### Enabling or Disabling User Private Attributes

- Step 1** Log in to FusionInsight Manager as user **admin** or a user bound to the Manager\_administrator role.
- Step 2** Choose **System > Permission > Security Policy > Independent Configurations**.
- Step 3** Toggle on or off **Independent Attribute**, enter the password as prompted, and click **OK**.

**Step 4** After the identity is authenticated, wait until the OMS configuration is modified and click **Finish**.

 **NOTE**

After the independent attribute is disabled:

- A user who has the attribute can cancel it from the drop-down list of the username in the upper right corner of the page. The user cannot set the independent attribute again once it is cancelled. After the attribute is cancelled, existing independent tables will retain the attribute. However, the user cannot create independent tables again.
- Users without this attribute cannot set or cancel the attribute.

----End

## Configuring the Independent Attribute

**Step 1** Log in to FusionInsight Manager as a service user.

---

**NOTICE**

Administrators cannot initialize the password of the user after the independent attribute is set. If the user password is forgotten, the password cannot be retrieved.

User **admin** cannot set the independent attribute.

---

**Step 2** Move the cursor to the username in the upper right corner of the page.

**Step 3** Select **Set Independent** or **Cancel Independent**.

 **NOTE**

- If the independent attribute is toggled on and has been set for the service user, **Cancel Independent** is displayed.
- If the independent attribute is toggled on but has been cancelled for the service user, **Set Independent** is displayed.
- If the independent attribute is toggled off but has been set for the service user, **Cancel Independent** is displayed.
- If the independent attribute is toggled off and has been cancelled for the service user, no option related to the independent attribute is displayed.

**Step 4** Enter the password as prompted and click **OK**.

**Step 5** After the identity is authenticated, click **OK** in the dialog box.

----End

## 6.11 Managing MRS Cluster Metadata

### 6.11.1 MRS Cluster Metadata Overview

MRS's data connections manage external source connections used by components in a cluster. For example, they associate Hive metadata with external relational databases.

- Local metadata: Metadata is stored in the local GaussDB of a cluster. When the cluster is deleted, the metadata is also deleted. To retain the metadata, manually back up the metadata in the database in advance.
- External data connection: You can associate the MRS cluster with an external data connection that is in the same VPC and subnet as the MRS cluster. Metadata is stored in the associated database and is not deleted even if the cluster is deleted. Multiple MRS clusters can share the same metadata.

MRS clusters support the following types of external data connections:

- Clusters with Hive installed can connect to RDS for PostgreSQL. The database version is **PostgreSQL14**.
- Clusters where Hive or Ranger is installed can connect to RDS for MySQL. The database version is **MySQL 5.7.x/MySQL 8.0**.
- Only MRS 3.1.2-LTS.3, MRS 3.1.5, and MRS 3.3.0-LTS clusters can connect to GaussDB(for MySQL).
- If the cluster supports LakeFormation, you can select LakeFormation as the data connection.

When Hive metadata is switched between different clusters, MRS synchronizes only the permissions in the metadata database of the Hive component. The permission model on MRS is maintained on MRS Manager. Therefore, when Hive metadata is switched between clusters, the permissions of users or user groups cannot be automatically synchronized to MRS Manager of another cluster.

## 6.11.2 Storing Ranger Metadata to RDS

This topic describes how to switch the Ranger metadata of the existing cluster to the metadata stored in the RDS database. This operation enables MRS clusters to share the same metadata, retains the metadata when the cluster is deleted, and avoids Ranger metadata migration during cluster migration.

### Disabling Ranger Authentication for Cluster Components

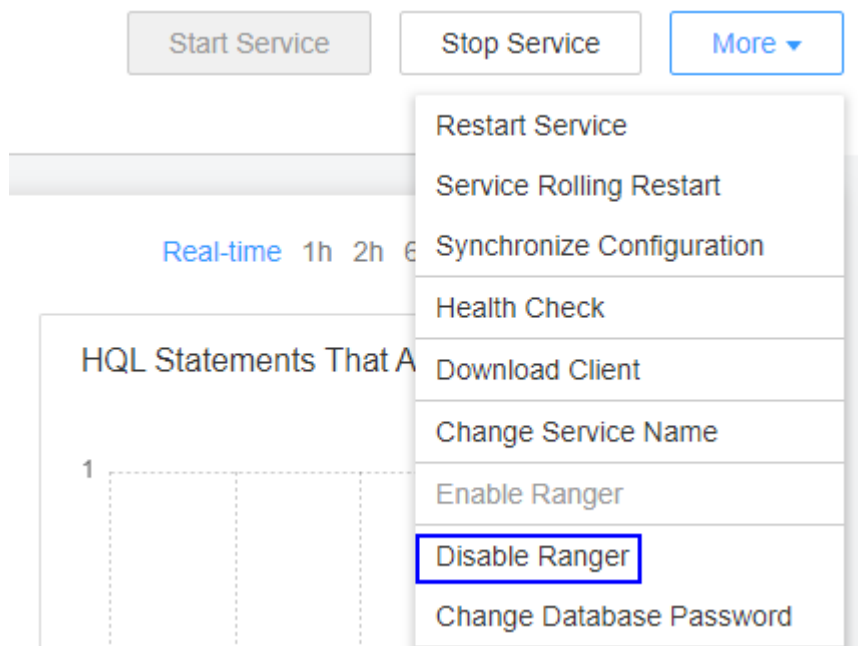
This operation is required only for **MRS 3.1.0 or later**.

**Step 1** Log in to FusionInsight Manager and choose **Cluster > Services > Service Name**.

Currently, the following components in an MRS 3.1.x cluster support Ranger authentication: HDFS, HBase, Hive, Spark, Impala, Storm, and Kafka.

**Step 2** In the upper right corner of the **Dashboard** page, click **More** and select **Disable Ranger**. If **Disable Ranger** is dimmed, Ranger authentication is disabled, as shown in [Figure 6-72](#).


**Figure 6-72** Disabling Ranger authentication



**Step 3** (Optional) To use an existing authentication policy, perform this step to export the authentication policy on the Ranger web page. After the Ranger metadata is switched, you can import the existing authentication policy again. The following uses Hive as an example. After the export, a policy file in JSON format is generated in a local directory.

1. Log in to FusionInsight Manager.
2. Choose **Cluster > Services > Ranger** to go to the Ranger service overview page.
3. Click **RangerAdmin** in the **Basic Information** area to go to the Ranger web UI.

The **admin** user in Ranger belongs to the **User** type. To view all management pages, click the username in the upper right corner and select **Log Out** to log out of the system.


4. Log in to the system as user **rangeradmin** (default password: **Rangeradmin@123**) or another user who has the Ranger administrator permissions. For details about the user and its default password, see [User Account List](#).
5. Click the export button  in the row where the Hive component is located to export the authentication policy.

**Figure 6-73** Exporting authentication policies



6. Click **Export**. After the export is complete, a policy file in JSON format is generated in a local directory.

**Figure 6-74** Exporting Hive authentication policies



----End

## Creating and Configuring an RDS DB Instance

**Step 1** Log in to the RDS console and buy an RDS DB instance. For details, see [Buying a DB Instance](#).

### NOTE


- To ensure network communications between the cluster and the MySQL or PostgreSQL database, create the instance in the same VPC and subnet as the cluster.
- Security group rules of the RDS DB instance must allow inbound access from MySQL (default port 3306) and PostgreSQL (default port 5432).

For example, click the instance name on the RDS console to go to the instance management page. In the **Connection Information** area, click the name next to **Security Group**. On the page that is displayed, click the **Inbound Rules** tab, and click **Add Rule**. In the displayed **Add Inbound Rule** dialog box, in the **Protocol & Port** area, select **TCP** and enter port number **3306**. In the **Source** area, select **IP address** and enter the IP addresses of all nodes where the MetaStore instances of Hive are located.

- Ranger can interconnect with RDS for MySQL databases of the **MySQL 5.7.x and 8.0** versions only.
- Hive can interconnect with RDS for MySQL and PostgreSQL databases. The supported versions are **MySQL 5.7.x and 8.0** and **PostgreSQL14**.

**Step 2** In the navigation pane of the RDS management console, choose **Instances**. Locate the row containing the RDS DB instance used by MRS data connections, click **More** in the **Operation** column, and select **Log In** to log in to the DB instance as user **root**.

**Figure 6-75** Logging in to an RDS DB instance



**Step 3** On the home page of the instance, click **Create Database** to create a database.

### NOTE

If no new database is created, the MRS data connections will fail to configure.



**Figure 6-76** Creating a database

The screenshot shows a 'Create Database' dialog box. The title bar contains the text 'Create Database' and a close button (X). The dialog has two input fields: 'Name' with a placeholder 'Database Name' and 'Character Set' with 'utf8' selected. A warning message 'Only user databases can be created' is displayed between the fields. At the bottom, there are 'OK' and 'Cancel' buttons.

**Step 4** On the top of the page, choose **Account Management > User Management**.

**NOTE**

- For clusters earlier than MRS 3.x, if the selected data connection is **RDS MySQL database**, ensure that the database user is **root**. If the user is not **root**, create a user and grant permissions to the user by referring to [Step 4](#) to [Step 6](#).
- For MRS 3.x or later clusters, when **Type** is set to **RDS MySQL database**, **Username** must not be **root**. In this case, create a user and grant permissions to the user by referring to [Step 4](#) to [Step 6](#).

**Step 5** Click **Create User** to create a non-root user and select all permissions listed in **Global Permissions**.

**NOTE**

If you are configuring an external RDS data connection for Ranger, you can select only the SELECT, INSERT, CREATE, RELOAD, CREATE USER, and GRANT permissions.

**Figure 6-77** Creating a user

The screenshot shows the 'Create User' form in the 'User Management' section. The form is organized into several sections:

- Basic Information:** Contains fields for Username (filled with 'mrs\_test01'), Host (filled with '%'), Password, and Confirm Password.
- Advanced Settings:** A section that is currently collapsed.
- Global Permissions:** A list of permissions with checkboxes:
  - Permission
  - SELECT
  - INSERT
  - UPDATE
  - DELETE
  - CREATE
- Object Permissions:** A section that is currently collapsed.
- Role:** A section that is currently collapsed.

**Step 6** On the top of the page, choose **SQL Operations > SQL Query**, switch to the target database by database name, and run the following SQL statements to grant permissions to the database user. In the following statements, *db\_name* and *db\_user* indicate the name of the database to be connected to MRS and the name of the new user, respectively.

```
grant all privileges on db_name.* to db_user'@%' with grant option;
grant reload on *.* to db_user'@%' with grant option;
flush privileges;
```

**Figure 6-78** Assigning permissions to database users



----End

## Creating an RDS Data Connection for an Existing MRS Cluster

Perform the following steps to create an RDS data connection for an existing MRS cluster.

- Step 1** Log in to the MRS management console, and choose **Data Connections** in the left navigation pane.
- Step 2** Click **Create Data Connection**.
- Step 3** Configure parameters according to [Table 6-61](#).

**Table 6-61** Parameters for creating a data connection

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type              | The type of an external source connection. Value options are as follows: <ul style="list-style-type: none"> <li>• <b>RDS PostgreSQL database.</b> Clusters with Hive installed can connect to this type of database.</li> <li>• <b>RDS MySQL database.</b> Clusters with Hive or Ranger installed can connect to this type of database.</li> </ul> |
| Name              | The name of a data connection.                                                                                                                                                                                                                                                                                                                     |
| Database Instance | The RDS database instance. This instance must be created in RDS before being referenced here, and the database must have been created. For details, see <a href="#">Creating and Configuring an RDS DB Instance</a> . Click <b>View DB Instance</b> to view the created DB instances.                                                              |
| Database          | The name of the database to be connected to.                                                                                                                                                                                                                                                                                                       |
| Username          | The username for logging in to the database to be connected.                                                                                                                                                                                                                                                                                       |
| Password          | The password for logging in to the database to be connected.                                                                                                                                                                                                                                                                                       |

 NOTE

If the selected data connection is an **RDS MySQL** database, ensure that the database user is a **root** user. If the user is not **root**, perform operations by referring to [Creating and Configuring an RDS DB Instance](#).

**Step 4** Click **OK**.

----End

## Configuring a Ranger Data Connection

**Step 1** Log in to the MRS console.

**Step 2** Click the name of the cluster to view its details.

**Step 3** Click **Manage** on the right of **Data Connection** to go to the data connection configuration page.

**Step 4** Click **Configure Data Connection** and set related parameters.

- **Component Name:** Ranger
- **Module Type:** Ranger metadata
- **Connection Type:** RDS MySQL database
- **Connection Instance:** Select a created RDS MySQL DB instance. For details about how to create a data connection, see [Creating an RDS Data Connection for an Existing MRS Cluster](#).

**Step 5** Select **I understand the consequences of performing the scale-in operation** and click **Test**.

**Step 6** After the test is successful, click **OK** to complete the data connection configuration.

**Step 7** Log in to FusionInsight Manager.

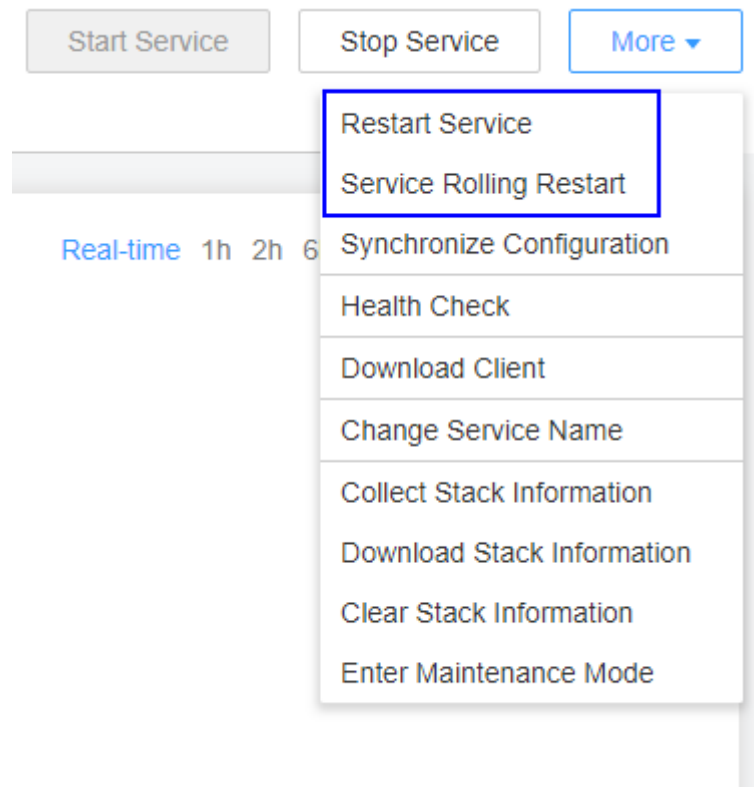
**Step 8** Choose **Cluster > Services > Ranger** to go to the Ranger service overview page.

**Step 9** Choose **More > Restart Service** or **More > Service Rolling Restart**.

If you choose **Restart Service**, services will be interrupted during the restart. If you select **Service Rolling Restart**, rolling restart can minimize the impact or do not affect service running.

Restarting Ranger will affect the permissions of all components controlled by Ranger and may affect service running. Restart Ranger when the cluster is idle or during off-peak hours. Before the Ranger component is restarted, the policies in the Ranger component still take effect.

**Figure 6-79** Restarting a service





**Step 10** Enable Ranger authentication for the component to be authenticated. The Hive component is used as an example.

Currently, the following components in an MRS 3.1.x cluster support Ranger authentication: HDFS, HBase, Hive, Spark, Impala, Storm, and Kafka.


1. Log in to FusionInsight Manager and choose **Cluster** > **Services** > *Service Name*.
2. In the upper right corner of the **Dashboard** page, click **More** and select **Enable Ranger**.

**Figure 6-80** Enabling Ranger authentication



**Step 11** Log in to the Ranger web UI and click the import button  in the row of the Hive component.

**Figure 6-81** Clicking the import button



**Step 12** Import parameters.

- Click **Select file** and select the authentication policy file downloaded in [Step 3.6](#).
- Select **Merge If Exist Policy**.

**Figure 6-82** Importing authentication policies

**Import Policy** ✕

**ⓘ** 'Override Policy' has higher priority than 'Merge If Exist Policy', if user selects both of them, then only 'Override Policy' take effect.

**Select File :**  
Select file  Merge If Exist Policy:  Override Policy:   
Ranger\_Policies\_20210331\_180915.json ✕

**ⓘ** All services gets listed on service destination when Zone destination is blank. When zone is selected at destination, then only services associated with that zone will be listed.

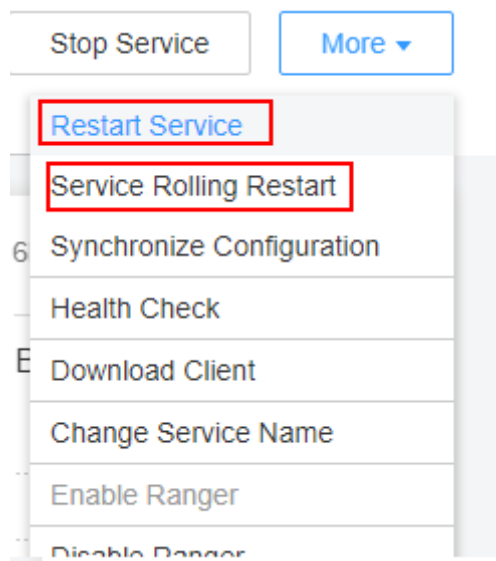
**Specify Zone Mapping :**  
Source:  To: Destination:

**Specify Service Mapping:**  
Source:  To: Destination:  ✕

**Step 13** Restart the component for which Ranger authentication is enabled.

1. Log in to FusionInsight Manager.
2. Choose **Cluster > Services > Hive** to go to the Hive service overview page.
3. Choose **More > Restart Service** or **More > Service Rolling Restart**.

**Figure 6-83** Restarting a service



If you choose **Restart Service**, services will be interrupted during the restart. If you select **Service Rolling Restart**, rolling restart can minimize the impact or do not affect service running.

----End

### 6.11.3 Storing Hive Metadata to RDS

This section describes how to switch the Hive metadata of an active cluster to the metadata stored in a local database or RDS database. This operation enables MRS clusters to share the same metadata, retains the metadata when the cluster is deleted, and avoids Hive metadata migration during cluster migration.

#### Creating and Configuring an RDS DB Instance

**Step 1** Log in to the RDS console and buy an RDS DB instance. For details, see [Buying a DB Instance](#).

##### NOTE


- To ensure network communications between the cluster and the MySQL or PostgreSQL database, create the instance in the same VPC and subnet as the cluster.
- Security group rules of the RDS DB instance must allow inbound access from MySQL (default port 3306) and PostgreSQL (default port 5432).

For example, click the instance name on the RDS console to go to the instance management page. In the **Connection Information** area, click the name next to **Security Group**. On the page that is displayed, click the **Inbound Rules** tab, and click **Add Rule**. In the displayed **Add Inbound Rule** dialog box, in the **Protocol & Port** area, select **TCP** and enter port number **3306**. In the **Source** area, select **IP address** and enter the IP addresses of all nodes where the MetaStore instances of Hive are located.

- Ranger can interconnect with RDS for MySQL databases of the **MySQL 5.7.x and 8.0** versions only.
- Hive can interconnect with RDS for MySQL and PostgreSQL databases. The supported versions are **MySQL 5.7.x and 8.0** and **PostgreSQL14**.

**Step 2** In the navigation pane of the RDS management console, choose **Instances**. Locate the row containing the RDS DB instance used by MRS data connections, click **More** in the **Operation** column, and select **Log In** to log in to the DB instance as user **root**.

**Figure 6-84** Logging in to an RDS DB instance



**Step 3** On the home page of the instance, click **Create Database** to create a database.

##### NOTE

If no new database is created, the MRS data connections will fail to configure.



**Figure 6-85** Creating a database

The screenshot shows a 'Create Database' dialog box. The title bar contains the text 'Create Database' and a close button (X). The dialog has two input fields: 'Name' with a placeholder 'Database Name' and 'Character Set' with 'utf8' selected. A warning message 'Only user databases can be created' is displayed below the Name field. At the bottom, there are 'OK' and 'Cancel' buttons.

**Step 4** On the top of the page, choose **Account Management > User Management**.

**NOTE**

- For clusters earlier than MRS 3.x, if the selected data connection is **RDS MySQL database**, ensure that the database user is **root**. If the user is not **root**, create a user and grant permissions to the user by referring to [Step 4](#) to [Step 6](#).
- For MRS 3.x or later clusters, when **Type** is set to **RDS MySQL database**, **Username** must not be **root**. In this case, create a user and grant permissions to the user by referring to [Step 4](#) to [Step 6](#).

**Step 5** Click **Create User** to create a non-root user and select all permissions listed in **Global Permissions**.

**NOTE**

If you are configuring an external RDS data connection for Ranger, you can select only the SELECT, INSERT, CREATE, RELOAD, CREATE USER, and GRANT permissions.

**Figure 6-86** Creating a user

The screenshot shows the 'Create User' form in the 'User Management' section. The form is organized into several sections:

- Basic Information:** Contains fields for Username (filled with 'mrs\_test01'), Host (filled with '%'), Password, and Confirm Password.
- Advanced Settings:** A section that is currently collapsed.
- Global Permissions:** A list of permissions with checkboxes:
  - Permission
  - SELECT
  - INSERT
  - UPDATE
  - DELETE
  - CREATE
- Object Permissions:** A section that is currently collapsed.
- Role:** A section that is currently collapsed.

**Step 6** On the top of the page, choose **SQL Operations > SQL Query**, switch to the target database by database name, and run the following SQL statements to grant permissions to the database user. In the following statements, *db\_name* and *db\_user* indicate the name of the database to be connected to MRS and the name of the new user, respectively.

```
grant all privileges on db_name.* to db_user'@%' with grant option;
grant reload on *.* to db_user'@%' with grant option;
flush privileges;
```

**Figure 6-87** Assigning permissions to database users



----End

## Creating an RDS Data Connection for an Existing MRS Cluster

Perform the following steps to create an RDS data connection for an existing MRS cluster.

- Step 1** Log in to the MRS management console, and choose **Data Connections** in the left navigation pane.
- Step 2** Click **Create Data Connection**.
- Step 3** Configure parameters according to [Table 6-62](#).

**Table 6-62** Parameters for creating a data connection

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type              | The type of an external source connection. Value options are as follows: <ul style="list-style-type: none"> <li>• <b>RDS PostgreSQL database.</b> Clusters with Hive installed can connect to this type of database.</li> <li>• <b>RDS MySQL database.</b> Clusters with Hive or Ranger installed can connect to this type of database.</li> </ul> |
| Name              | The name of a data connection.                                                                                                                                                                                                                                                                                                                     |
| Database Instance | The RDS database instance. This instance must be created in RDS before being referenced here, and the database must have been created. For details, see <a href="#">Creating and Configuring an RDS DB Instance</a> . Click <b>View DB Instance</b> to view the created DB instances.                                                              |
| Database          | The name of the database to be connected to.                                                                                                                                                                                                                                                                                                       |
| Username          | The username for logging in to the database to be connected.                                                                                                                                                                                                                                                                                       |
| Password          | The password for logging in to the database to be connected.                                                                                                                                                                                                                                                                                       |

 NOTE

If the selected data connection is an **RDS MySQL** database, ensure that the database user is a **root** user. If the user is not **root**, perform operations by referring to [Creating and Configuring an RDS DB Instance](#).

**Step 4** Click **OK**.

----End

## Configuring a Hive Data Connection

This function is not supported in MRS 3.0.5.

**Step 1** Log in to the MRS console. In the navigation pane on the left, choose **Active Clusters**.

**Step 2** Click the name of a cluster to go to the cluster details page.

**Step 3** On the **Dashboard** tab page, click **Manage** next to **Data Connection**.

**Step 4** On the **Data Connections** page, the data connections of the cluster are displayed. You can click **Disassociate** to delete a data connection.

**Step 5** If there is no associated data connection on the **Data Connection** dialog box, click **Configure Data Connection** to add a connection.

 NOTE

Only one data connection can be configured for a module type. For example, after a data connection is configured for Hive metadata, no other data connection can be configured for it. If no module type is available, the **Configure Data Connection** button is unavailable.

**Table 6-63** Configuring a Hive data connection

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Component            | Hive                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Module Type          | Hive metadata                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Data Connection Type | <ul style="list-style-type: none"><li>• RDS PostgreSQL database (supported for clusters of MRS 1.9.x)</li><li>• RDS MySQL database</li><li>• Local database</li></ul>                                                                                                                                                                                                                                                                                     |
| Instance             | This parameter is valid only when <b>Data Connection Type</b> is set to <b>RDS PostgreSQL database</b> or <b>RDS MySQL database</b> . Select the name of the connection between the MRS cluster and the RDS database. This instance must be created before being referenced here. You can click <b>Create Data Connection</b> to create a data connection. For details, see <a href="#">Creating an RDS Data Connection for an Existing MRS Cluster</a> . |

**Figure 6-88** Configuring a data connection

**Configure Data Connection**

Component Name

Module Type

Data Connection Type

Instance  [Create Data Connection](#)

**Step 6** Click **Test** to test connectivity of the data connection.

**Step 7** After the data connection is successful, click **OK**.

**Step 8** Upload the open-source driver package of PostgreSQL or MySQL to all nodes where MetaStore instances are deployed to replace the existing driver package of the cluster.

- PostgreSQL: Upload the driver package **postgresql-42.2.5.jar** to the **\$ {BIGDATA\_HOME}/third\_lib/Hive** directory on all MetaStore instance nodes. Download the open-source driver package from <https://repo1.maven.org/maven2/org/postgresql/postgresql/42.2.5/>.
- MySQL: Log in to the MySQL official website (<https://www.mysql.com/>), choose **Downloads > Community > MySQL Connectors > Connector/J** to download the driver package of the required version, and upload the driver package to the **/opt/Bigdata/FusionInsight\_HD\_\*/install/FusionInsight-Hive-\*/hive-\*/lib/** directory on all Metastore instance nodes.

**Step 9** On the cluster details page, click the **Components** tab and click **Hive**. On the service details page, choose **More > Restart Service** and click **OK** to restart the Hive service.

**NOTE**

- If IAM users are not synchronized, click the **Dashboard** tab on the cluster details page, click **Synchronize** on the right of **IAM User Sync**, and then restart the Hive service.
- Hive will create necessary database tables in the specified database after restart. (If tables already exist, they will not be created.)

----End

## 6.11.4 Configuring a LakeFormation Data Connection

### 6.11.4.1 LakeFormation Overview

LakeFormation is a one-stop enterprise-class data lake and warehouse construction service. It provides APIs and a GUI for unified management of data lake metadata, and is compatible with Hive metadata and Ranger permission models. LakeFormation can connect to multiple compute engines and big data cloud services seamlessly to ensure quick building and easy operation of data lakes and unleash rich value of service data.

You can create a LakeFormation instance and connect it to an MRS cluster for centrally manage data lake metadata and permission.


### Restrictions and Constraints

- **Before interconnecting LakeFormation with MRS clusters, pay attention to the following restrictions:**
  - MRS clusters and LakeFormation instances must belong to the same cloud account and region.
  - The VPC of the access client created by LakeFormation must be in the same VPC of the MRS clusters.
  - The MRS cluster can only interconnect with the catalog named **hive** in the LakeFormation instance.
  - For existing MRS clusters, you need to migrate the metadata database and permission policies to the LakeFormation instance, and then configure the interconnection.
  - If metadata in multiple MRS clusters needs to be migrated to the same LakeFormation instance, the database names of the MRS clusters must be unique.
- **After MRS is interconnected with LakeFormation, MRS components are subject to the following constraints:**
  - Hive does not support temporary tables.
  - Hive does not support cross-cluster column encryption.
  - Hive WebHCat cannot interconnect with LakeFormation.
  - Hive cannot create an internal table if the designated directory already contains files.
  - Before creating a Hudi table, you need to add the path authorization of the Hudi table directory on LakeFormation to grant OBS read and write permissions.
  - Fields in a Hudi table cannot be edited on the LakeFormation console. You can only add, delete, or modify table fields on the Hudi client.
  - When Flink reads and writes Hive tables, only **hive\_sync.mode=jdbc** can be used to synchronize Hive tables. HMS is not supported.
  - If a low-permission user lacks OBS path access permission for the default database, Spark will display a permission error message but will still successfully create the database.
- **After MRS is interconnected with LakeFormation, the permission policy restrictions are as follows:**


- In LakeFormation authorization, only LakeFormation roles can be used as authorization entities. IAM users or user groups cannot be used as authorization entities.
- The PolicySync process does not modify the default policy of the RangerAdmin Hive module in the cluster. The default policy still takes effect.
- After the PolicySync process is started, it compares the permissions with those of LakeFormation instances and deletes the non-default policies that do not exist in LakeFormation. You are advised to migrate the permission policies to LakeFormation instances first.
- For the Hive module on the RangerAdmin web UI, do not add or delete non-default policies. Grant permissions on the data permission page of LakeFormation instances.
- After the interconnection between the MRS cluster and LakeFormation is canceled, the non-default policies of RangerAdmin will not be cleared. You need to manually clear them.
- Hive does not support SQL statements for granting permissions. You need to grant permissions on the **Data Permissions** page.
- MRS does not support LakeFormation row filtering.

## 6.11.4.2 Preparing for a LakeFormation Data Connection

### Creating a LakeFormation Instance

- Step 1** Log in to the Huawei Cloud management console, click  in the upper left corner, and choose **Analytics > DataArts LakeFormation**. The LakeFormation console is displayed.
- Step 2** Click **Buy Instance** in the upper right corner of the page and create a LakeFormation instance by referring to [Creating a LakeFormation Instance](#).
- Step 3** Create the **hive** catalog and the **default** database. If the catalog and database already exist in the instance, skip this step. For details, see [Managing Metadata](#).

MRS can interconnect with LakeFormation only when the catalog name of the LakeFormation instance is **hive**.

1. Ensure that the instance name displayed in the upper left corner is the name of the newly created instance. Then, choose **Metadata > Catalog** page. You only need one **hive** catalog per instance. If you already have one in the current instance, you can skip creating another one.
2. Click **Create** in the upper right corner. On the displayed page, set the following parameters, and click **Submit**.
  - **Catalog Name:** Enter **hive**. Do not use custom names.
  - **Select Location:** Click  to select the OBS storage path corresponding to the catalog, for example, **obs://lakeformation-test/hive** (which must be an existing path), and click **OK**.
  - Set other parameters based on the site requirements.

< Create Catalog

---

**Basic Information**

\* Catalog Name  Select  +

Location

Catalog Type  Description

---

**Database Storage Locations**

| No. | Location |
|-----|----------|
|     |          |

---

- In the navigation pane on the left, choose **Metadata > Database**, click **Create Database**, configure the following parameters, and click **Submit**. You only need one **default** database per instance. If you already have one in the current instance, you can skip creating another one.
  - Database Name:** Enter **default**. Do not use custom names.
  - Catalog:** **hive**
  - Select Location:** Click **+** to select a location in the Hive Catalog storage path, for example, **obs://lakeformation-test/hive/default** (which must be an existing path), and click **OK**.
  - Set other parameters based on the site requirements.

< Create Database

---

**Basic Information**

\* Database  \* Catalog

Name

\* Select  + Description

Location

---

**Data Table Storage Locations**

| No. | Location |
|-----|----------|
|     |          |

---

**Function Storage Locations**

| No. | Location |
|-----|----------|
|     |          |

---


**Step 4** Choose **Data Permissions > Data Authorization** in the navigation pane on the left. On the displayed page, you can grant permissions to access Hive Catalog to users and user groups based on service requirements. For details, see [Granting Permissions](#).

**Step 5** Choose **Clients** in the navigation pane on the left. Click **Create** to create a client for access management. The VPC and Subnet must be the same as those of the MRS cluster you want to interconnect. For details, see [Managing Access Clients](#).



×



### Create Client

If no VPC or subnet is available, [create one](#) .

\* Client

\* VPC

\* Subnet

 Ensure that your current account has enough quotas for resources like VPCEP and private DNS. Otherwise, the creation will fail, the client will be rolled back and deleted, and any related resources will be reclaimed. 

You can obtain the VPC and subnet of the MRS cluster on the **Dashboard** page of the cluster on the MRS management console.

Go to the client details page, record the access IP address of the client.

----End

## Creating an Agency for Interconnecting with LakeFormation

**Step 1** Log in to Huawei Cloud management console and go to the **Identity and Access Management** console page.

**Step 2** In the navigation pane on the left, choose **Agencies**. Click **Create Agency** in the upper right corner, set related parameters, and click **Next**.

Refer to the following descriptions to configure the parameters:

- **Agency Name:** For example, enter **visit\_lakeformation\_agency**.
- **Agency Type:** Select **Account**.
- **Delegated Account:** Enter the name of the delegated Huawei Cloud account.
- **Validity Period:** Set this parameter as you need.

Figure 6-89 Creating an agency

\* Agency Name

\* Agency Type  Account  
 Delegate another Huawei Cloud account to perform operations on your resources.  
 Cloud service  
 Delegate a cloud service to access your resources in other cloud services.

\* Delegated Account

\* Validity Period

Description

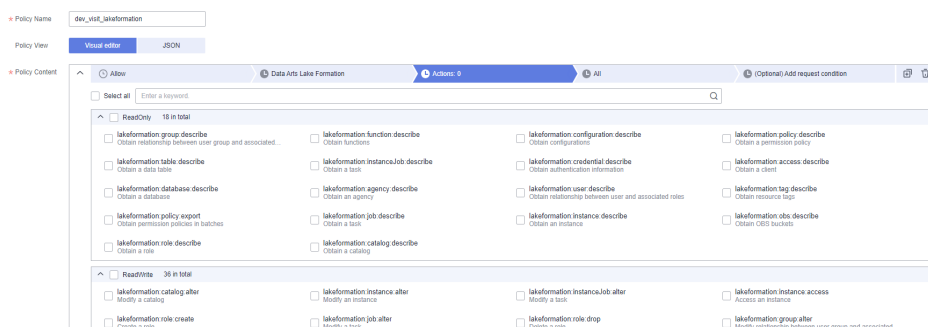
0/255

**Step 3** In the upper right corner of the **Select Policy/Role** page, click **Create Policy**, configure the following information, and click **Next**.

- **Policy Name:** For example, enter **dev\_visit\_lakeformation**.
- **Policy View:** Select **Visual editor** or **JSON**.
- **Policy Content:**

The policy must contain **lakeformation:policy:export** and **lakeformation:role:describe**. Set other parameters as you need.

- If you are using the visual editor, set the **Policy Content** as follows: Select **Data Arts LakeFormation** for **Cloud Service** and select the required permissions in **Actions**. Set other parameters as you need.



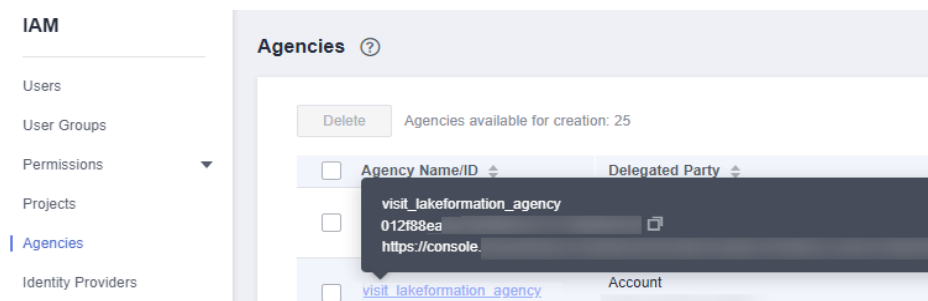
- If you are using the JSON view, configure the policy content by referring to the following content:

```
{
 "Version": "1.1",
 "Statement": [
 {
 "Effect": "Allow",
```

```
"Action": [
 "lakeformation:table:create",
 "lakeformation:database:alter",
 "lakeformation:table:alter",
 "lakeformation:database:drop",
 "lakeformation:database:create",
 "lakeformation:role:describe",
 "lakeformation:policy:create",
 "lakeformation:policy:export",
 "lakeformation:function:alter",
 "lakeformation:function:describe",
 "lakeformation:table:drop",
 "lakeformation:catalog:describe",
 "lakeformation:table:describe",
 "lakeformation:function:drop",
 "lakeformation:database:describe",
 "lakeformation:function:create",
 "lakeformation:transaction:operate",
 "lakeformation:policy:drop"
]
}
]
```

- Step 4** Select the new policy name, for example, **dev\_visit\_lakeformation**, and click **Next**.
- Step 5** Select the desired scope requiring minimum authorization and click **OK** to create an agency.
- Step 6** On the **Agencies** page, move the cursor to the name of the newly created agency to obtain the ID of the agency that has the permission to access LakeFormation.

**Figure 6-90** Viewing an agency ID



----End

## Creating an Agency for Interconnecting with OBS

- Step 1** Log in to Huawei Cloud management console and go to the **Identity and Access Management** console page.
- Step 2** In the navigation pane on the left, choose **Agencies**. Click **Create Agency** in the upper right corner, set related parameters, and click **Next**.

Refer to the following descriptions to set the parameters:

- **Agency Name:** For example, enter **visit\_obs\_agency**.
- **Agency Type:** Select **Account**.
- **Delegated Account:** Enter the name of the delegated Huawei Cloud account.

- **Validity Period:** Set this parameter as you need.

**Step 3** In the upper right corner of the **Select Policy/Role** page, click **Create Policy**, configure the following information, and click **Next**.

- **Policy Name:** For example, enter **dev\_visit\_obs**.
- **Policy View:** Select **JSON**.
- **Policy Content:** Enter the following information.

```
{
 "Version": "1.1",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "obs:bucket:GetBucketLocation",
 "obs:bucket:ListBucketMultipartUploads",
 "obs:object:GetObject",
 "obs:object:ModifyObjectMetaData",
 "obs:object:DeleteObject",
 "obs:object:ListMultipartUploadParts",
 "obs:bucket:HeadBucket",
 "obs:object:AbortMultipartUpload",
 "obs:bucket:ListBucket",
 "obs:object:PutObject"
],
 "Resource": [
 "OBS:*:*:bucket:*",
 "OBS:*:*:object:*"
]
 }
]
}
```

#### NOTE

In the **Resource** parameter, **bucket** indicates the OBS bucket name, and **object** indicates the OBS object name. Specify the names as needed. If this parameter is set to \*, the policy applies to all OBS buckets or objects.

- Set other parameters as you need.

**Step 4** Select the new policy name, for example, **dev\_visit\_obs**, and click **Next**.

**Step 5** Select the desired scope requiring minimum authorization and click **OK** to create an agency.

**Step 6** On the **Agencies** page, move the cursor to the name of the newly created agency to obtain the ID of the agency that has the permission to access OBS.

----End

## Creating an ECS/BMS Agency

**Step 1** Log in to Huawei Cloud management console and go to the **Identity and Access Management** console page.

**Step 2** In the navigation pane on the left, choose **Agencies**. Click **Create Agency** in the upper right corner, set related parameters, and click **Next**.

Refer to the following descriptions to set the parameters:

- **Agency Name:** For example, enter **lakeformation\_test**.
- **Agency Type:** Select **Cloud service**.

- **Cloud Service:** Select **ECS BMS**.
- **Validity Period:** Set this parameter as you need.

**Step 3** In the upper right corner of the **Select Policy/Role** page, click **Create Policy**, configure the following information, and click **Next**.

- **Policy Name:** Enter a policy name.
- **Policy View:** Select **JSON**.
- **Policy Content:** Configure the parameter as follows.

```
{
 "Version": "1.1",
 "Statement": [
 {
 "Action": [
 "iam:agencies:assume"
],
 "Resource": {
 "uri": [
 "/iam/agencies/ID of the agency that grants LakeFormation access permission to your account",
 "/iam/agencies/ID of the agency that grants the OBS access permission to your account"
]
 },
 "Effect": "Allow"
 }
]
}
```

#### NOTE

- To obtain the ID of the agency that has the permission to access LakeFormation, refer to [Step 6](#).
- To obtain the ID of the agency that has the permission to access OBS, refer to [Step 6](#).

**Step 4** Select the name of the newly created custom agency and click **Next**.

**Step 5** Select the desired scope requiring minimum authorization and click **OK** to create an agency.

----End

## Creating a LakeFormation Connection

#### NOTE

You need to contact technical support to be added in the whitelist before you create a LakeFormation data connection.

**Step 1** Log in to the MRS management console, and choose **Data Connections** in the left navigation pane.

**Step 2** Click **Create Data Connection**.

**Step 3** Set parameters by referring to [Table 6-64](#) and click **OK**.

**Table 6-64** LakeFormation data connection

| Parameter              | Description                                                                                                                                                                                                                                                                  |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type                   | Select <b>LakeFormation</b> . Only MRS 3.3.0-LTS supports this connection type.                                                                                                                                                                                              |
| Name                   | The name of a data connection.                                                                                                                                                                                                                                               |
| LakeFormation Instance | Select an instance name.<br>This instance is created in the LakeFormation console before being referenced here. For details about how to create the instance, see <a href="#">Creating a LakeFormation Instance</a> . Click <b>View instances</b> to view created instances. |
| VPC                    | It must be in the same VPC as the MRS cluster you want to interconnect.                                                                                                                                                                                                      |
| Subnet                 | Subnet name                                                                                                                                                                                                                                                                  |
| VPC Endpoint           | Select a VPC endpoint or click <b>Create VPC endpoints</b> for LakeFormation Instance.<br>After you select a VPC endpoint, you will be charged by the VPCEP service.                                                                                                         |
| LakeFormation Agency   | Select <b>Available agencies</b> and select the agency created in <a href="#">Creating an Agency for Interconnecting with LakeFormation</a> , for example, <code>visit_lakeformation_agency</code> .                                                                         |

**Figure 6-91** Creating a LakeFormation data connection

**Create Data Connection** ×

Type: LakeFormation ↻

Name:

LakeFormation Instance:  ↻

[View instances](#)

VPC:  ↻ ? [View VPC](#)

Subnet:  ↻ ? [View Subnet](#)

VPC Endpoint:  ↻ ? [View VPC endpoints](#)

[Create VPC endpoint](#)

LakeFormation Agency: MRS\_LAKEFORMA... Available agencies ?

↻

[Manage external source connections used by components in this cluster. Learn more](#)

OK Cancel

**Step 4** Record the ID of the created data connection on the **Data Connections** page.

----End

## Obtaining the Account ID

**Step 1** Log in to the management console as the user for interconnecting MRS with LakeFormation.

**Step 2** Hover over the username and choose **My Credentials** from the drop-down list.

**Step 3** On the **API Credentials** page, obtain the **Account ID** and view the **Project ID** in the project list.

IAM User Name:  Account Name:

IAM User ID:  Account ID: 20490i

---

Projects

| Project ID | Project Name | Region |
|------------|--------------|--------|
| 1c0fbdb    |              |        |

**Step 4** Grant the current user the permission to use LakeFormation.

1. Click in the upper left corner and choose **Analytics > LakeFormation**.
2. Check whether the service authorization page is displayed, or go to the Service Authorization page to check whether the service has been authorized.

- Select the checkbox to agree with the LakeFormation Service Statement and click **Authorize**.
- If no, the current user has the permission to perform operations on LakeFormation.

----End

### 6.11.4.3 Configuring a LakeFormation Data Connection During Cluster Creation

You can configure a LakeFormation data connection when you create an MRS 3.3.0-LTS cluster. You need to configure MRS cluster parameters to interconnect with LakeFormation after the cluster is created.



#### Configuring a LakeFormation Data Connection During Cluster Creation



- Step 1** Go to the [Buy Cluster](#) page.
- Step 2** Click **Buy Cluster**. The page for buying a cluster is displayed.
- Step 3** On the page for buying a cluster, click the **Custom Config** tab.
- Step 4** Create a cluster by referring to [Configuring Custom Topology](#). The cluster must meet the requirements described in [Table 6-65](#).

**Table 6-65** Data connection parameters

| Parameter       | Description                                                                                                                                                                       |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Version Type    | LTS                                                                                                                                                                               |
| Cluster Version | Version of the MRS cluster to be interconnected<br>Currently, LakeFormation data connections can be configured during creation of MRS 3.3.0-LTS clusters and later versions only. |




| Parameter                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                        |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|----------------------|--------------------------------------------|--------------|--------------------------------------------------------------------------------------------|-------------------------------------------|---------------|-------------------------------------------------------------------------------------------------------|--------------------------------|--------|----------------------------------------------------------|------------------------------------------|-------|------------------------------------------------------------------------------------------------------------------------|---------------------------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|-----------------|------------------------------------------------------------------------------------|--------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|--------|----------------------------------------------------------------------------------------|--------------------------------|-------|-------------------------------|-----------------------------------------------|-------|------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|-------|-----------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------|-------|-------------------------------------------------------------------------------------------------------------|------------------------------|--------|----------------------------------------------------------------------------------------------------------|-------------------------------------|----------|-----------------------------------------------------------------------------------------------------------------------|--------------------------------|-------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|-------|------------------------------------------------------------------------------|----------------------------------------------|-------|---------------------------------------------------------------------------|
| Component                                      | <p>Components such as Hadoop, Ranger, Hive, Guardian, Spark (optional), and Flink (optional) must be included. For example, the following figure shows an example configuration. The configuration may vary depending on the cluster version.</p>  <p>The screenshot shows a table of components with the following columns: Name, Version, and Description. The components listed are:</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Version</th> <th>Description</th> </tr> </thead> <tbody> <tr><td><input checked="" type="checkbox"/> Hadoop</td><td>3.3.1</td><td>A framework that allows for the distributed processing of large data sets across clusters.</td></tr> <tr><td><input checked="" type="checkbox"/> Spark</td><td>3.3.1</td><td>Apache Spark is a fast and general engine based on open source Spark for large-scale data processing.</td></tr> <tr><td><input type="checkbox"/> HBase</td><td>2.4.14</td><td>HBase - distributed, versioned, non-relational database.</td></tr> <tr><td><input checked="" type="checkbox"/> Hive</td><td>3.1.0</td><td>Data warehouse software that facilitates query and management of large datasets stored in distributed storage systems.</td></tr> <tr><td><input type="checkbox"/> Loader</td><td>1.99.3</td><td>Loader is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured databases such as relational databases.</td></tr> <tr><td><input type="checkbox"/> Kafka</td><td>2.152.8.1</td><td>Apache Kafka is publish-subscribe messaging rethought as a distributed commit log.</td></tr> <tr><td><input type="checkbox"/> Flume</td><td>1.11.0</td><td>Flume is a distributed, reliable, and available service for efficiently collecting, aggregating, and moving large amounts of log data.</td></tr> <tr><td><input checked="" type="checkbox"/> Flink</td><td>1.15.0</td><td>Apache Flink is an open source platform for scalable batch and stream data processing.</td></tr> <tr><td><input type="checkbox"/> Oozie</td><td>5.1.0</td><td>Hadoop job scheduling system.</td></tr> <tr><td><input checked="" type="checkbox"/> ZooKeeper</td><td>3.8.1</td><td>A centralized service for maintaining configuration information, naming, performing distributed synchronization, and providing group services.</td></tr> <tr><td><input type="checkbox"/> HdfsEngine</td><td>2.0.0</td><td>HdfsEngine is a distributed SQL query engine designed to query large data sets distributed over one or more heterogeneous data sources.</td></tr> <tr><td><input checked="" type="checkbox"/> Ranger</td><td>2.3.0</td><td>RANGER is a framework to enable, monitor and manage comprehensive data security across the Hadoop platform.</td></tr> <tr><td><input type="checkbox"/> Tez</td><td>0.10.2</td><td>An application framework which allows for a complex directed acyclic graph of tasks for processing data.</td></tr> <tr><td><input type="checkbox"/> ClickHouse</td><td>23.3.237</td><td>ClickHouse is a column-oriented database management system (DBMS) for online analytical processing of queries (OLAP).</td></tr> <tr><td><input type="checkbox"/> HlTDB</td><td>1.1.0</td><td>Apache HlTDB (Database for Internet of Things) is an IoT native database with high performance for data management and analysis, deployment on the edge and the cloud.</td></tr> <tr><td><input checked="" type="checkbox"/> JobGateway</td><td>1.0.0</td><td>JobGateway service provides the job submission capability through REST APIs.</td></tr> <tr><td><input checked="" type="checkbox"/> Guardian</td><td>0.1.0</td><td>Guardian provides temporary Authentication credentials for accessing OBS.</td></tr> </tbody> </table> | Name                                                                                                                                                                   | Version                  | Description          | <input checked="" type="checkbox"/> Hadoop | 3.3.1        | A framework that allows for the distributed processing of large data sets across clusters. | <input checked="" type="checkbox"/> Spark | 3.3.1         | Apache Spark is a fast and general engine based on open source Spark for large-scale data processing. | <input type="checkbox"/> HBase | 2.4.14 | HBase - distributed, versioned, non-relational database. | <input checked="" type="checkbox"/> Hive | 3.1.0 | Data warehouse software that facilitates query and management of large datasets stored in distributed storage systems. | <input type="checkbox"/> Loader | 1.99.3             | Loader is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured databases such as relational databases. | <input type="checkbox"/> Kafka | 2.152.8.1       | Apache Kafka is publish-subscribe messaging rethought as a distributed commit log. | <input type="checkbox"/> Flume | 1.11.0         | Flume is a distributed, reliable, and available service for efficiently collecting, aggregating, and moving large amounts of log data. | <input checked="" type="checkbox"/> Flink | 1.15.0 | Apache Flink is an open source platform for scalable batch and stream data processing. | <input type="checkbox"/> Oozie | 5.1.0 | Hadoop job scheduling system. | <input checked="" type="checkbox"/> ZooKeeper | 3.8.1 | A centralized service for maintaining configuration information, naming, performing distributed synchronization, and providing group services. | <input type="checkbox"/> HdfsEngine | 2.0.0 | HdfsEngine is a distributed SQL query engine designed to query large data sets distributed over one or more heterogeneous data sources. | <input checked="" type="checkbox"/> Ranger | 2.3.0 | RANGER is a framework to enable, monitor and manage comprehensive data security across the Hadoop platform. | <input type="checkbox"/> Tez | 0.10.2 | An application framework which allows for a complex directed acyclic graph of tasks for processing data. | <input type="checkbox"/> ClickHouse | 23.3.237 | ClickHouse is a column-oriented database management system (DBMS) for online analytical processing of queries (OLAP). | <input type="checkbox"/> HlTDB | 1.1.0 | Apache HlTDB (Database for Internet of Things) is an IoT native database with high performance for data management and analysis, deployment on the edge and the cloud. | <input checked="" type="checkbox"/> JobGateway | 1.0.0 | JobGateway service provides the job submission capability through REST APIs. | <input checked="" type="checkbox"/> Guardian | 0.1.0 | Guardian provides temporary Authentication credentials for accessing OBS. |
| Name                                           | Version                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Description                                                                                                                                                            |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input checked="" type="checkbox"/> Hadoop     | 3.3.1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | A framework that allows for the distributed processing of large data sets across clusters.                                                                             |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input checked="" type="checkbox"/> Spark      | 3.3.1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Apache Spark is a fast and general engine based on open source Spark for large-scale data processing.                                                                  |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> HBase                 | 2.4.14                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | HBase - distributed, versioned, non-relational database.                                                                                                               |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input checked="" type="checkbox"/> Hive       | 3.1.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Data warehouse software that facilitates query and management of large datasets stored in distributed storage systems.                                                 |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> Loader                | 1.99.3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Loader is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured databases such as relational databases.                          |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> Kafka                 | 2.152.8.1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Apache Kafka is publish-subscribe messaging rethought as a distributed commit log.                                                                                     |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> Flume                 | 1.11.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Flume is a distributed, reliable, and available service for efficiently collecting, aggregating, and moving large amounts of log data.                                 |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input checked="" type="checkbox"/> Flink      | 1.15.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Apache Flink is an open source platform for scalable batch and stream data processing.                                                                                 |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> Oozie                 | 5.1.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Hadoop job scheduling system.                                                                                                                                          |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input checked="" type="checkbox"/> ZooKeeper  | 3.8.1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | A centralized service for maintaining configuration information, naming, performing distributed synchronization, and providing group services.                         |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> HdfsEngine            | 2.0.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | HdfsEngine is a distributed SQL query engine designed to query large data sets distributed over one or more heterogeneous data sources.                                |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input checked="" type="checkbox"/> Ranger     | 2.3.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | RANGER is a framework to enable, monitor and manage comprehensive data security across the Hadoop platform.                                                            |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> Tez                   | 0.10.2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | An application framework which allows for a complex directed acyclic graph of tasks for processing data.                                                               |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> ClickHouse            | 23.3.237                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | ClickHouse is a column-oriented database management system (DBMS) for online analytical processing of queries (OLAP).                                                  |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input type="checkbox"/> HlTDB                 | 1.1.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Apache HlTDB (Database for Internet of Things) is an IoT native database with high performance for data management and analysis, deployment on the edge and the cloud. |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input checked="" type="checkbox"/> JobGateway | 1.0.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | JobGateway service provides the job submission capability through REST APIs.                                                                                           |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| <input checked="" type="checkbox"/> Guardian   | 0.1.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Guardian provides temporary Authentication credentials for accessing OBS.                                                                                              |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| Metadata                                       | <p>Select <b>External data connection</b> and set the following parameters:</p> <ol style="list-style-type: none"> <li><b>LakeFormation Metadata:</b> Enable this function.</li> <li><b>LakeFormation Instance:</b> Select the LakeFormation data connection created in <a href="#">Creating a LakeFormation Connection</a>.</li> <li><b>Data Connection Type:</b> Retain the default value.</li> </ol> <p>For example, the following figure shows an example configuration. The configuration may vary depending on the cluster version.</p>  <p>The screenshot shows the Metadata configuration page with the following settings:</p> <ul style="list-style-type: none"> <li>Local: <input type="radio"/> External data connection: <input checked="" type="radio"/></li> <li>LakeFormation Metadata: <input checked="" type="checkbox"/></li> <li>LakeFormation instance: [Dropdown menu]</li> <li>Create Data Connection: [Link]</li> </ul> <p>The table below shows the configuration for various data connections:</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Data Type</th> <th>Data Connection Type</th> <th>Data Connection Instance</th> <th>Instance VPC</th> </tr> </thead> <tbody> <tr> <td>Spark</td> <td>Metastore metadata</td> <td>LakeFormation</td> <td>[Dropdown menu]</td> <td>[Dropdown menu]</td> </tr> <tr> <td>Hive</td> <td>Hive metadata</td> <td>Local database</td> <td>--</td> <td>--</td> </tr> <tr> <td>Hive</td> <td>Metastore metadata</td> <td>LakeFormation</td> <td>[Dropdown menu]</td> <td>[Dropdown menu]</td> </tr> <tr> <td>Ranger</td> <td>Ranger metadata</td> <td>Local database</td> <td>--</td> <td>--</td> </tr> </tbody> </table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Name                                                                                                                                                                   | Data Type                | Data Connection Type | Data Connection Instance                   | Instance VPC | Spark                                                                                      | Metastore metadata                        | LakeFormation | [Dropdown menu]                                                                                       | [Dropdown menu]                | Hive   | Hive metadata                                            | Local database                           | --    | --                                                                                                                     | Hive                            | Metastore metadata | LakeFormation                                                                                                                                 | [Dropdown menu]                | [Dropdown menu] | Ranger                                                                             | Ranger metadata                | Local database | --                                                                                                                                     | --                                        |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| Name                                           | Data Type                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Data Connection Type                                                                                                                                                   | Data Connection Instance | Instance VPC         |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| Spark                                          | Metastore metadata                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | LakeFormation                                                                                                                                                          | [Dropdown menu]          | [Dropdown menu]      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| Hive                                           | Hive metadata                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Local database                                                                                                                                                         | --                       | --                   |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| Hive                                           | Metastore metadata                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | LakeFormation                                                                                                                                                          | [Dropdown menu]          | [Dropdown menu]      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| Ranger                                         | Ranger metadata                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Local database                                                                                                                                                         | --                       | --                   |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| VPC                                            | VPC where the LakeFormation data connection belongs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                        |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |
| Subnet                                         | Subnet name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                        |                          |                      |                                            |              |                                                                                            |                                           |               |                                                                                                       |                                |        |                                                          |                                          |       |                                                                                                                        |                                 |                    |                                                                                                                                               |                                |                 |                                                                                    |                                |                |                                                                                                                                        |                                           |        |                                                                                        |                                |       |                               |                                               |       |                                                                                                                                                |                                     |       |                                                                                                                                         |                                            |       |                                                                                                             |                              |        |                                                                                                          |                                     |          |                                                                                                                       |                                |       |                                                                                                                                                                        |                                                |       |                                                                              |                                              |       |                                                                           |


| Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Topology Adjustment     | <p>Cluster node topology. Select <b>Enable</b> and ensure that at least one PolicySync (PSC) instance is on a Ranger node which also has the RangerAdmin (RA) instance and at least two TokenSever (TS) instances are added to the Guardian component.</p> <p>For example, the following figure shows an example configuration. The configuration may vary depending on the cluster version.</p>  |
| Kerberos Authentication | <p>Enable this function.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Agency                  | <p>Select <b>Configure</b> next to <b>Set Advanced Options</b>, select <b>Available agencies</b> for <b>Agency</b>, and select the agency created in <a href="#">Creating an ECS/BMS Agency</a>.</p> <p>For example, the following figure shows an example configuration. The configuration may vary depending on the cluster version.</p>                                                      |

**Step 5** Click the name of the created MRS cluster on the **Active Clusters** page. In the displayed tab, click **Synchronize** next to **IAM User Sync** to synchronize IAM users.

### O&M Management

MRS Manager 

Access Manager 

**IAM User Sync** 

Synchronized **Synchronize**

**Step 6** Refer to [Configuring the MRS 3.3.0-LTS Cluster](#) to configure decoupled storage and compute and download the client.

----End

## Configuring the MRS 3.3.0-LTS Cluster

**Step 1** Log in to FusionInsight Manager of the MRS cluster. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).

**Step 2** Configure the Guardian component.

1. On FusionInsight Manager, click **Cluster**, choose **Services > Guardian**, click **Configurations**, and then **All Configurations**. Search for and modify the following parameters, and click **Save**.

**Table 6-66** Guardian parameters

| Parameter                             | Description                                                                                                                                                                                                                                                                    | Value                                                                              |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| token.server.access.iam.domain.id     | Account ID of the user who accesses IAM<br>Obtain the account ID by referring to <a href="#">Obtaining the Account ID</a> .                                                                                                                                                    | xxx                                                                                |
| token.server.access.iam.project.id    | Project ID of the user who accesses IAM<br>Obtain the project ID by referring to <a href="#">Obtaining the Account ID</a> .                                                                                                                                                    | xxx                                                                                |
| token.server.access.label.agency.name | Name of an IAM agency.<br>The agency must have the permission to access OBS.<br>Set the value to the name of the agency created in <a href="#">Creating an Agency for Interconnecting with OBS</a> .                                                                           | visit_obs_agency                                                                   |
| fs.obs.delegation.token.providers     | Name of the class that generates <b>delegation.token</b> . The default value is empty.<br>Select both the following parameters: <ul style="list-style-type: none"><li>- com.huawei.mrs.dt.MRSDelegationTokenProvider</li><li>- com.huawei.mrs.dt.GuardianDTPProvider</li></ul> | com.huawei.mrs.dt.MRSDelegationTokenProvider,com.huawei.mrs.dt.GuardianDTPProvider |
| fs.obs.guardian.accesslabel.enabled   | Whether to enable <b>access label</b> for using Guardian to connect to OBS.                                                                                                                                                                                                    | true                                                                               |

| Parameter               | Description                 | Value |
|-------------------------|-----------------------------|-------|
| fs.obs.guardian.enabled | Whether to enable Guardian. | true  |

- In the **Dashboard** tab of the Guardian service, choose **More > Restart Service**.


**Step 3** Interconnecting Hive with OBS

- Log in to FusionInsight Manager, click **Cluster**, choose **Services > Hive**, click **Configurations**, and then **All Configurations**.
- In the navigation pane on the left, choose **HiveServer > Customization**. Configure the following parameters.

**Table 6-67** Custom HiveServer parameters

| Parameter                      | Description                                                                                                                                                                                                                                               | Example Value                                                                                                                      |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| hive.server.customized.configs | <ul style="list-style-type: none"> <li>Add the <b>hive.metastore.warehouse.dir</b> parameter.</li> <li>Set this parameter to the OBS storage path of the hive Catalog, which is obtained in <a href="#">Creating a LakeFormation Instance</a>.</li> </ul> | <ul style="list-style-type: none"> <li>Name: hive.metastore.warehouse.dir</li> <li>Value: obs://lakeformation-test/hive</li> </ul> |

**Figure 6-92** Configuring **hive.metastore.warehouse.dir**



- Click **Save**.

**Step 4** Interconnecting Spark with OBS If your cluster does not have the Spark component, skip this step.

- Log in to FusionInsight Manager and choose **Cluster > Services > Spark**. Click **Configurations** then **All Configurations**.
- In the navigation pane on the left, choose **JDBCServer > Customization**. Add custom parameters and set them by referring to the following table.

**Table 6-68** Spark parameters

| Custom Item                        | Value                                                                                                                                                                                                                                                                                                    |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| spark.hive-site.customized.configs | <ul style="list-style-type: none"><li>Parameter: <b>hive.metastore.warehouse.dir</b></li><li>Value: Set this parameter to the OBS storage path of the hive Catalog, which is obtained in <a href="#">Creating a LakeFormation Instance</a>. For example, <b>obs://lakeformation-test/hive</b>.</li></ul> |

- In the navigation pane on the left, choose **SparkResource** > **Customization**. Set parameters by referring to [Table 6-68](#).
- Click **Save**.

**Step 5** In the **Components** tab of the MRS cluster, check whether there are expired components. If there are, click **Restart** in the **Operation** column to restart these components.

**Step 6** Download and reinstall the complete MRS cluster client. For details, see [Installing a Client \(MRS 3.x or Later\)](#).

**Step 7** Update the built-in client configuration file of the cluster to submit jobs on the management console.

On the dashboard page of the MRS cluster, obtain the EIP, use this IP address to log in to a Master node, and run the following commands to refresh the built-in client of the cluster:

```
su - omm
```

```
sh /opt/executor/bin/refresh-client-config.sh
```

**Step 8** Log in to the node where the client is installed and check the database on the Hive client to verify that the interconnection is successful.

```
source Client installation path//bigdata_env
```

```
kinit Component service user
```

```
beeline
```

```
show databases;desc database default;
```

```
!q
```

**Step 9** On the Spark client, check the database to ensure that the interconnection is successful. If your cluster does not have the Spark component, skip this step.

```
source Client installation path//Spark/component_env
```

```
spark-sql
```

```
show databases;desc database default;
```

```
----End
```

## 6.11.5 Managing MRS Cluster Data Connections

This topic describes how to create, view, and delete a cluster data connection on the MRS console.

### Creating a Data Connection

**Step 1** Log in to the MRS management console. In the navigation pane, choose **Data Connections**.

**Step 2** Click **Create Data Connection**.

For details about how to configure an RDS for MySQL data connection, see [Creating an RDS Data Connection for an Existing MRS Cluster](#).

For details about how to configure a LakeFormation data connection, see [Creating a LakeFormation Connection](#).

**Step 3** Click **OK**.

----End

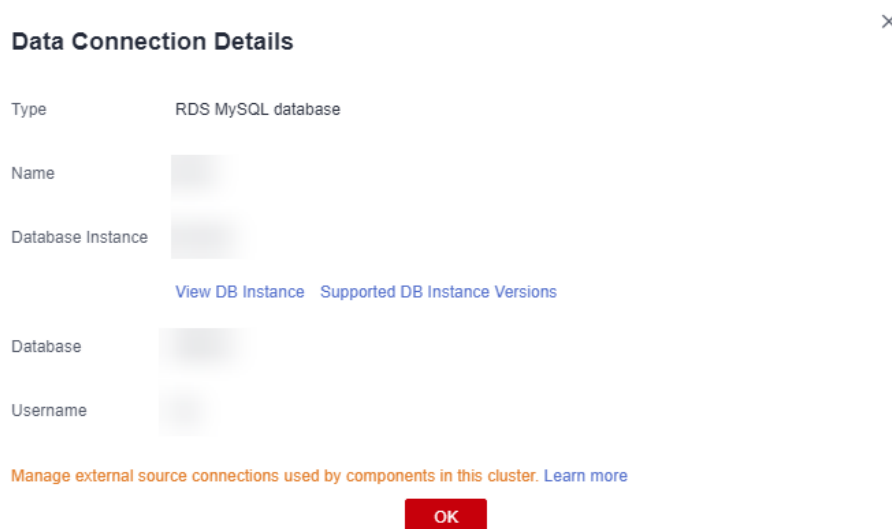
### Checking Data Connection Information

**Step 1** Log in to the MRS management console. In the navigation pane, choose **Data Connections**.

**Step 2** In the data connection list, click the desired data connection. On the page that is displayed, view its details.

For example, the data connection information of the RDS MySQL database is as follows:

**Figure 6-93** Viewing the data connection information of the RDS MySQL database



----End

## Deleting a Data Connection

- Step 1** Log in to the MRS management console. In the navigation pane, choose **Data Connections**.
- Step 2** In the **Operation** column of the data connection list, click **Delete** in the row where the data connection to be deleted is located. Enter **DELETE** in the **Delete Data Connection** dialog box and click **OK**.

If the selected data connection has been associated with a cluster, the deletion does not affect the cluster.

----End

# 6.12 Managing Static Service Resources in an MRS Cluster

## 6.12.1 Overview of Static Service Resources

### Introduction

A cluster allocates static service resources to services Flume, HBase, HDFS, and YARN. The total volume of computing resources allocated to each service is fixed, and they are static. A tenant can exclusively use or share a service to obtain the resources required for running this service.

### Static Service Pool

Static service pools are used to specify service resource configurations.

Static service pools centrally manage resources that can be used by each service.

- Limits the total number of resources that can be used by each service. Specifically, the total number of CPU, I/O, and memory resources can be configured on the nodes where services Flume, HBase, HDFSioTDB, Kafka (Kafka supports static service pools only in MRS 3.2.0 or later), and YARN are deployed.
- Isolates the resources of services in a cluster from those of other services. In this way, the load of one service has very limited impact on other services.

### Scheduling Mechanism

The time-based dynamic resource scheduling mechanism enables different volumes of static resources to be configured for services at different time, optimizing service running environments and improving the cluster efficiency.

In a complex cluster environment, multiple services share resources in the cluster, but the resource service period of each service may be different.

The following use a bank customer as an example:

- The HBase query service is heavy in the daytime.

- The query service is light, but the Hive analysis service is heavy at night.

If fixed resources are allocated to each service, the following problems may occur:

- The query service cannot obtain sufficient resources while the resources for the analysis service are idle in the daytime.
- The analysis service cannot obtain sufficient resources while the resources for the query service are idle at night.

As a result, the cluster resource utilization is low and the service capability is weak. Resolve the problem in the following ways:

- Sufficient resources need to be configured for HBase in the daytime.
- Sufficient resources need to be configured for Hive at night.

The time-based dynamic scheduling mechanism can efficiently utilize resources and run tasks.

## 6.12.2 Configuring Static Resources for an MRS Cluster

You can adjust resource base on Manager and customize resource configuration groups if you need to control service resources used on each node in a cluster or the available CPU or I/O quotas on each node at different time segments.

### Impact on the System

- After a static service pool is configured, the configuration of affected services expires. You need to restart the services. Services are unavailable during restart.
- After a static service pool is configured, the maximum number of resources used by each service and role instance cannot exceed the upper limit.

### Configuring Static Resources (MRS 3.x and Later)

#### Modify the Resource Adjustment Base

**Step 1** Log in to FusionInsight Manager and choose **Cluster > Static Service Pool Configurations**.

**Step 2** Click **Configuration** in the upper right corner. The page for configuring resource pools is displayed.

**Step 3** Change the values of **CPU (%)** and **Memory (%)** in the **System Resource Adjustment Base** area.

Modifying the system resource adjustment base changes the maximum physical CPU and memory usage on nodes by services. If multiple services are deployed on the same node, the maximum physical resource usage of all services cannot exceed the adjusted CPU or memory usage.

**Step 4** Click **Next**.

To modify parameters again, click **Previous**.

#### Modify the Default Resource Configuration Group

**Step 5** Click **default**. In the **Configure weight** table, set **CPU LIMIT(%)**, **CPU SHARE(%)**, **I/O(%)**, and **Memory(%)** for each service.



**Figure 6-94** Weight Configuration

| Services | CPU LIMIT (%) | CPU SHARE (%) | I/O (%) | Memory (%) |
|----------|---------------|---------------|---------|------------|
| Flume    | 0             | 0             | 0       | 0          |
| HBase    | 0             | 0             | 0       | 0          |
| HDFS     | 0             | 0             | 0       | 0          |
| Impala   | 0             | 0             | 0       | 0          |
| Kudu     | 0             | 0             | 0       | 0          |
| Yarn     | 0             | 0             | 0       | 0          |
| Total:   | 0             | 0             | 0       | 0          |

**NOTE**

- The sum of **CPU LIMIT(%)** and **CPU SHARE(%)** used by all services can exceed 100%.
- The sum of **I/O(%)** used by all services can exceed 100% but cannot be 0.
- The sum of **Memory(%)** used by all services can be greater than, smaller than, or equal to 100%.
- **Memory(%)** cannot take effect dynamically and can only be modified in the default configuration group.
- **CPU LIMIT(%)** is used to configure the ratio of the number of CPU cores that can be used by a service to those can be allocated to related nodes.
- **CPU SHARE(%)** is used to configure the ratio of the time when a service uses a CPU core to the time when other services use the CPU core. That is, the ratio of time when multiple services compete for the same CPU core.

**Step 6** Click **Generate detailed configurations based on weight configurations**. FusionInsight Manager generates the actual values of the parameters in the default weight configuration table based on the cluster hardware resources and allocation information.

**Step 7** Click **OK**.

In the displayed dialog box, click **OK**.

**Add a Customized Resource Configuration Group**

**Step 8** Determine whether to automatically adjust resource configurations at different time segments.

- If yes, go to [Step 9](#).
- If no, use the default configurations, and no further action is required.

**Step 9** Click **Configuration**, change the system resource adjustment base values, and click **Next**.

**Step 10** Click **Add** to add a resource configuration group.

**Figure 6-95** Adding a resource configuration group

The screenshot shows a web interface for configuring resource groups. At the top, there are tabs for 'default', 'dynamic-config1', and 'Add'. Below this, the interface is divided into three steps:

- Step 1: Scheduling Time** - A link labeled 'Configuration' is visible.
- Step 2: Weight Configuration** - This step contains a table with columns for 'Services', 'CPU LIMIT (%)', 'CPU SHARE (%)', and 'I/O (%)'. The table lists services: Flume, HBase, HDFS, Impala, Kudu, and Yarn, each with input fields set to '0'. A 'Total' row shows 0 for all columns. A 'Service Name' search box is on the right.
- Step 3:** A button labeled 'Generate detailed configurations based on weight configurations' and another 'Service Name' search box.

**Step 11** In **Step 1: Scheduling Time**, click **Configuration**.

The page for configuring the time policy is displayed.

Modify the following parameters based on service requirements and click **OK**.

- **Repeat:** If this parameter is selected, the customized resource configuration is applied repeatedly based on the scheduling period. If this parameter is not selected, set the date and time when the configuration of the group of resources can be applied.
- **Repeat Policy:** The available values are **Daily**, **Weekly**, and **Monthly**. This parameter is valid only when **Repeat** is selected.
- **On:** indicates the time period between the start time and end time when the resource configuration is applied. Set a unique time range. If the time range overlaps with that of an existing group of resource configuration, the time range cannot be saved.

 NOTE

- The default group of resource configuration takes effect in all undefined time segments.
- The newly added resource group is a parameter set that takes effect dynamically in a specified time range.
- The newly added resource group can be deleted. A maximum of four resource configuration groups that take effect dynamically can be added.
- Select a repetition policy. If the end time is earlier than the start time, the resource configuration ends in the next day by default. For example, if a validity period ranges from 22:00 to 06:00, the customized resource configuration takes effect from 22:00 on the current day to 06:00 on the next day.
- If the repetition policy types of multiple configuration groups are different, the time ranges can overlap. The policy types are listed as follows by priority from low to high: daily, weekly, and monthly. The following is an example. There are two resource configuration groups using the monthly and daily policies, respectively. Their application time ranges in a day overlap as follows: 04:00 to 07:00 and 06:00 to 08:00. In this case, the configuration of the group that uses the monthly policy prevails.
- If the repetition policy types of multiple resource configuration groups are the same, the time ranges of different dates can overlap. For example, if there are two weekly scheduling groups, you can set the same time range on different day for them, such as to 04:00 to 07:00, on Monday and Wednesday, respectively.

**Step 12** Modify the resource configuration of each service in **Step 2: Weight Configuration**.

**Step 13** Click **Generate detailed configuration**. FusionInsight Manager generates the actual values of the parameters in the default weight configuration table based on the cluster hardware resources and allocation information.

**Step 14** Click **OK**.

In the displayed dialog box, click **OK**.

----End

## Configuring Static Resources (MRS 2.x and Earlier)

**Step 1** Modify the system resource adjustment base.

1. On MRS Manager, click **System**. In the **Resource** area, click **Configure Static Service Pool**.
2. Click **Configuration**. The service pool configuration group management page is displayed.
3. In the **System Resource Adjustment Base** area, change the values of **CPU(%)** and **Memory(%)**.

Modifying **System Resource Adjustment Base** limits the maximum physical CPU and memory resource percentage of nodes that can be used by the Flume, HBase, HDFS, Impala and YARN services. If multiple services are deployed on the same node, the maximum physical resource usage of all services cannot exceed the adjusted CPU or memory usage.

4. Click **Next**.

If you need to modify the parameters again, click **Previous** in the lower part of the page.

**Step 2** Modify the **default** configuration group of the service pool.

1. In the **Service Pool Configuration** table, set **CPU LIMIT(%)**, **CPU SHARE(%)**, **I/O(%)**, and **Memory(%)** for the Flume, HBase, HDFS, Impala, and YARN services.

 **NOTE**

- The sum of **CPU LIMIT(%)** used by all services can exceed 100%.
  - The sum of **CPU SHARE(%)** and **I/O(%)** used by all services must be 100%. For example, if CPU resources are allocated to the HDFS and YARN services, the total CPU resources allocated to the two services are 100%.
  - The sum of **Memory(%)** used by all services can be greater than, smaller than, or equal to 100%.
  - **Memory(%)** cannot take effect dynamically and can only be modified in the default configuration group.
2. Click in the blank area of the page to complete the editing. MRS Manager generates the correct values of service pool parameters in the **Detailed Configuration** area based on the cluster hardware resources and allocation information.
  3. You can click the edit icon on the right of **Detailed Configuration** to modify the parameter values of the service pool based on service requirements.


In the **Service Pool Configuration** area, click the specified service name. The **Detailed Configuration** area displays only the parameters of the service. Manually changing parameter values does not refresh the service resource usage. In added configuration groups, the configuration group numbers of the parameters that take effect dynamically will be displayed. For example, HBase: RegionServer: dynamic-config1.RES\_CPUSET\_PERCENTAGE. The parameter functions do not change.

**Table 6-69** Parameters of the static service pool

| Parameter                                                          | Description                                         |
|--------------------------------------------------------------------|-----------------------------------------------------|
| - RES_CPUSET_PERCENTAGE<br>- dynamic-configX.RES_CPUSET_PERCENTAGE | Configures the service CPU percentage.              |
| - RES_CPU_SHARE<br>- dynamic-configX.RES_CPU_SHARE                 | Configures the service CPU usage share.             |
| - RES_BLKIO_WEIGHT<br>- dynamic-configX.RES_BLKIO_WEIGHT           | Configures service I/O usage.                       |
| HBASE_HEAPSIZE                                                     | Configures the maximum JVM memory for RegionServer. |
| HADOOP_HEAPSIZE                                                    | Configures the maximum JVM memory of a DataNode.    |

| Parameter                           | Description                                                                     |
|-------------------------------------|---------------------------------------------------------------------------------|
| yarn.nodemanager.resource.memory-mb | Configures the memory that can be used by NodeManager on the current node.      |
| dfs.datanode.max.locked.memory      | Configures the maximum memory that can be used by a DataNode as the HDFS cache. |
| FLUME_HEAPSIZE                      | Configures the maximum JVM memory that can be used by each Flume instance.      |
| IMPALAD_MEM_LIMIT                   | Configures the maximum memory that can be used by an Impalad instance.          |

**Step 3** Add a customized resource configuration group.

1. Determine whether to automatically adjust resource configurations based on the time.  
If yes, go to [Step 3.2](#).  
If no, go to [Step 4](#).
2. Click  to add a resource configuration group. In the **Scheduling Time** area, click the edit icon. The time policy configuration page is displayed.  
Modify the following parameters based on service requirements and click **OK**.
  - **Repeat**: If selected, the resource configuration group runs repeatedly based on the scheduling period. If not selected, set the date and time when the configuration of the group of resources can be applied.
  - **Repeat Policy**: can be set to **Daily**, **Weekly**, and **Monthly**. This parameter is valid only when **Repeat** is selected.
  - **Between**: indicates the time period between the start time and end time when the resource configuration is applied. Set a unique time range. If the time range overlaps with that of an existing group of resource configuration, the time range cannot be saved. This parameter is valid only when **Repeat** is selected.

 NOTE

- The default group of resource configuration takes effect in all undefined time segments.
  - The newly added resource group is a parameter set that takes effect dynamically in a specified time range.
  - The newly added resource group can be deleted. A maximum of four resource configuration groups that take effect dynamically can be added.
  - Select a repetition policy. If the end time is earlier than the start time, the resource configuration ends in the next day by default. For example, if a validity period ranges from 22:00 to 06:00, the customized resource configuration takes effect from 22:00 on the current day to 06:00 on the next day.
  - If the repetition policy types of multiple configuration groups are different, the time ranges can overlap. The policy types are listed as follows by priority from low to high: daily, weekly, and monthly. The following is an example. There are two resource configuration groups using the monthly and daily policies, respectively. Their application time ranges in a day overlap as follows: 04:00 to 07:00 and 06:00 to 08:00. In this case, the configuration of the group that uses the monthly policy prevails.
  - If the repetition policy types of multiple resource configuration groups are the same, the time ranges of different dates can overlap. For example, if there are two weekly scheduling groups, you can set the same time range on different day for them, such as to 04:00 to 07:00, on Monday and Wednesday, respectively.
3. On the **Service Pool Configuration** page, modify the resource configuration of each service. Click the blank area on the page to complete the editing, and go to [Step 4](#).

You can click the edit icon on the right of **Service Pool Configuration** to modify the parameters. Click the edit icon in the **Detailed Configuration** area to manually update the parameter values generated by the system based on service requirements.

**Step 4** Saves the settings.

Click **Save**. In the **Save Configuration** dialog box, select **Restart the affected services or instances**. Click **OK** to save the settings and restart related services.

After the system displays "Operation succeeded", click **Finish**. The service is started successfully.

----End

### 6.12.3 Checking the Static Resources of an MRS Cluster

The big data management platform can manage and isolate service resources that are not running on YARN using static service resource pools. The system supports time-based automatic adjustment of static service resource pools. This enables the cluster to automatically adjust the parameter values at different periods to ensure more efficient resource utilization.

System administrators can view the monitoring indicators of resources used by each service in the static service pool on Manager. The monitoring indicators are as follows:

- CPU usage of services
- Total disk I/O read rate of services

- Total disk I/O write rate of services
- Total used memory of services

## Checking Static Resources (MRS 3.x and Later)

**Step 1** Log in to FusionInsight Manager and choose **Cluster > Static Service Pool Configurations**.

**Step 2** In the configuration group list, click a configuration group, for example, **default**.

**Step 3** Check the system resource adjustment base values.

- **System Resource Adjustment Base** indicates the maximum volume of resources that can be used by each node in the cluster. If a node has only one service, the service exclusively occupies the available resources on the node. If a node has multiple services, all services share the available resources on the node.
- **CPU** indicates the maximum number of CPUs that can be used by services on a node.
- **Memory** indicates the maximum memory that can be used by services on a node.

**Step 4** In **Chart**, view the metric data of the cluster service resource usage.

### NOTE

- You can click **Add Service to Chart** to add static service resource data of specific services (up to 12 services) to the chart.
- For details about how to manage a chart, see [Viewing MRS Cluster Resource Monitoring Metrics](#).

----End

## Checking Static Resources (MRS 2.x and Earlier)

**Step 1** On MRS Manager, click **System**. In the **Resource** area, click **Configure Static Service Pool**.

**Step 2** Click **Status**.

**Step 3** Check the system resource adjustment base values.

- **System Resource Adjustment Base** indicates the maximum volume of resources that can be used by each node in the cluster. If a node has only one service, the service exclusively occupies the available resources on the node. If a node has multiple services, all services share the available resources on the node.
- **CPU(%)** indicates the maximum number of CPUs that can be used by services on a node.
- **Memory(%)** indicates the maximum memory that can be used by services on a node.

**Step 4** Check the cluster service resource usage.

In the chart area, select **All services** from the service drop-down list box. The resource usage status of all services in the service pool is displayed.

 **NOTE**

**Effective Configuration Group** indicates the resource control configuration group used by the cluster service. By default, the **default** configuration group is used at all time every day, indicating that the cluster service can use all CPUs and 70% memory of the node.

**Step 5** View the resource usage of a single service.

In the chart area, select a service from the service drop-down list box. The resource usage status of the service is displayed.

**Step 6** You can set the interval for automatically refreshing the page.

**Step 7** In the **Period** area, select a time range for viewing service resources.

**Step 8** Click **View** to view the service resource data in the corresponding time range.

**Step 9** Customize a service resource report.

1. Click **Customize** and select the service source indicators to be displayed.
2. Click **OK** to save the selected monitoring metrics for display.

 **NOTE**

Click **Clear** to cancel all the selected monitoring metrics in a batch.

**Step 10** Export a monitoring report.

Click **Export**. MRS Manager will generate a report about the selected service resources in a specified time of period. Save the report.

 **NOTE**

To view the curve charts of monitoring metrics in a specified period, click **View**.

----**End**



# 7 MRS Cluster O&M

---

## 7.1 Cluster O&M

### Account Maintenance Suggestions

It is recommended that the administrator conduct routine checks on the accounts. The check covers the following items:

- Check whether the accounts of the OS, Manager, and each component are necessary and whether temporary accounts have been deleted.
- Check whether the permissions of the accounts are appropriate. Different administrators have different rights.
- Check and audit the logins and operation records of all types of accounts.

### Password Maintenance Suggestions

Accessing portal requires identity authentication. The complexity and validity period of an account password must meet your security requirements.

Refer to the following suggestions to maintain passwords:

- Assign dedicated personnel to keep OS passwords.
- Use passwords that meet certain strength requirements, such as minimum password length or mixing of letter cases.
- Encrypt passwords before transferring them, and do not transfer them via email.
- Encrypt passwords for storage.
- Remind enterprise users to change passwords during system handover.
- Change passwords periodically.

### Log Maintenance Suggestions

Operation logs help discover exceptions such as illegal operations and login by unauthorized users. The system records important operations in logs. You can use operation logs to locate problems.

- **Checking Logs Regularly**  
Check system logs periodically and handle exceptions such as unauthorized operations or logins in a timely manner.
- **Backing Up Logs Regularly**  
Audit logs provided by Manager and clusters record user activity and operation information. You can export audit logs from Manager. If there are too many audit logs in the system, you can configure dump parameters to dump audit logs to a specified server to ensure that the cluster nodes disk space is sufficient.
- **Maintenance Owner**  
Network monitoring engineers and system maintenance engineers

## Manager Routine Maintenance

To ensure long-term and stable running of the system, administrators or maintenance engineers need to periodically check items listed in the following table and rectify the detected faults based on the check results. It is recommended that administrators or engineers record the result in each task scenario and sign off based on the enterprise management regulations.

**Table 7-1** Routine maintenance check items

| Routine Maintenance Frequency | Role                              | Check Item                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Daily                         | Check the cluster service status. | <ul style="list-style-type: none"> <li>• Check whether the running status and configuration status of each service are normal and whether the status icons are green.</li> <li>• Check whether the running status and configuration status of the role instances in each service are normal and whether the status icons are green.</li> <li>• Check whether the active/standby status of role instances in each service can be properly displayed.</li> <li>• Check whether the dashboard of the services and role instances can be displayed properly.</li> </ul> |
|                               | Check the cluster host status.    | <ul style="list-style-type: none"> <li>• Check whether the running status of each host is normal and whether the status icon is green.</li> <li>• Check the current disk usage, memory usage, and CPU usage of each host. Check whether the current memory usage and CPU usage are increasing.</li> </ul>                                                                                                                                                                                                                                                           |

| Routine Maintenance Frequency | Role                                 | Check Item                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                               | Check the cluster alarm information. | Check whether alarms were generated for unhandled exceptions on the previous day, including alarms that were automatically cleared.                                                                                                                                                                                                                                                                                      |
|                               | Check the cluster audit information. | Check whether critical and major operations are performed on the previous day and whether the operations are valid.                                                                                                                                                                                                                                                                                                      |
|                               | Check the cluster backup status.     | Check whether OMS, LDAP, DBService, and NameNode have been automatically backed up on the previous day.                                                                                                                                                                                                                                                                                                                  |
|                               | View the health check result.        | Perform a health check on Manager and download the health check report to check whether the current cluster is abnormal. You are advised to enable the automatic health check, export the latest cluster health check result, and repair unhealthy items based on the result.                                                                                                                                            |
|                               | Check the network communication.     | Check the cluster network status and check whether the network communication between nodes is delayed.                                                                                                                                                                                                                                                                                                                   |
|                               | Check the storage status.            | <p>Check whether the total data storage volume of the cluster increases abruptly.</p> <ul style="list-style-type: none"> <li>● Check whether the disk usage is close to the threshold. If yes, locate the causes. For example, check whether the junk data or cold data left by services needs to be cleared.</li> <li>● Check whether disk partitions need to be expanded based on the service growth trend.</li> </ul> |

| Route<br>Maintenance<br>Frequency | Role                           | Check Item                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                   | Check logs.                    | <ul style="list-style-type: none"> <li>• Check whether there are failed or unresponsive MapReduce and Spark tasks. Check the <code>/tmp/logs/\${username}/logs/\${application id}</code> log file in HDFS and rectify faults.</li> <li>• Check Yarn task logs, view the logs of failed and unresponsive tasks, and delete duplicate data.</li> <li>• Check the worker logs of Storm.</li> <li>• Back up logs to the storage server.</li> </ul> |
| Weekly                            | User management                | Check whether the user password is about to expire and notify the user of changing the password. To change the password of a machine-machine user, you need to download the keytab file again.                                                                                                                                                                                                                                                 |
|                                   | Analyze alarms.                | Export and analyze alarms generated in a specified period.                                                                                                                                                                                                                                                                                                                                                                                     |
|                                   | Scan disks.                    | Check the disk health status. You are advised to use a dedicated disk check tool.                                                                                                                                                                                                                                                                                                                                                              |
|                                   | Collect statistics on storage. | Check in batches whether the disk data of cluster nodes is evenly stored, filter out the disks whose data increases significantly or is insufficient, and check whether the disks are normal.                                                                                                                                                                                                                                                  |
|                                   | Record changes.                | Arrange and record the operations on cluster configuration parameters and files to provide reference for fault analysis and handling.                                                                                                                                                                                                                                                                                                          |
| Monthly                           | Analyze logs.                  | <ul style="list-style-type: none"> <li>• Collect and analyze hardware logs of cluster node servers, such as BMC system logs.</li> <li>• Collect and analyze the OS logs of the cluster node servers.</li> <li>• Collect and analyze cluster logs.</li> </ul>                                                                                                                                                                                   |
|                                   | Diagnose the network.          | Analyze the network health status of the cluster.                                                                                                                                                                                                                                                                                                                                                                                              |
|                                   | Manage hardware.               | Check the equipment room environment and clean the devices.                                                                                                                                                                                                                                                                                                                                                                                    |

## 7.2 Logging In to an MRS Cluster

### 7.2.1 Checking MRS Active/Standby Management Nodes

Some O&M operation scripts and commands need to be run or can be run only on the active management node. You can log in to a Master node or the Manager (MRS 3.x or later) to determine the active and standby management nodes (active and standby OMS nodes).

An active/standby switchover can be implemented between Master1 and Master2. For this reason, Master1 may not be the active management node.

#### Running the Script to Determine Active and Standby Nodes

**Step 1** Obtain information about the master nodes of the MRS cluster.

1. Log in to the MRS management console, choose **Active Clusters** and click the name of the target cluster to access its details page.
2. In the **Nodes** tab, view Master node names. The node that contains **master1** in its name is the Master1 node. The node that contains **master2** in its name is the Master2 node.

**Step 2** Determine the active and standby management nodes of the cluster.

1. Remotely log in to the Master1 node. For details, see [Logging In to an MRS Cluster Node](#).

Master nodes support Cloud-Init. The preset username for Cloud-Init is **root** and the password is the one you set during cluster creation.

2. Run the following commands to switch the user:

```
sudo su - root
```

```
su - omm
```

3. Run the following command to identify the active and standby management nodes:

- For versions earlier than **MRS 3.x**, run the following command:

```
sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh
```

- For **MRS 3.x** and later versions, run the following command:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh
```

In the command output, the node whose **HAActive** is **active** is the active management node (mgtomsdat-sh-3-01-1 in the following example), and the node whose **HAActive** is **standby** is the standby management node (mgtomsdat-sh-3-01-2 in the following example).

```
Ha mode
double
NodeName HostName HAVersion StartTime HAActive
HAAllResOK HARunPhase
192-168-0-30 mgtomsdat-sh-3-01-1 V100R001C01 20xx-11-18 23:43:02
active normal Activated
192-168-0-24 mgtomsdat-sh-3-01-2 V100R001C01 20xx-11-21 07:14:02
standby normal Deactivated
```

**NOTE**

If the Master1 node to which you have logged in is the standby management node and you need to log in to the active management node, run the following command:

```
ssh IP address of Master2 node
```

----End

## Logging in to Manager to Determine Active and Standby Nodes

This section applies only to MRS 3.x or later.

**Step 1** Log in to Manager. For details, see [Accessing MRS Manager](#).

**Step 2** Click **Hosts**. The **Hosts** page is displayed.

**Step 3** View and record the IP addresses of the active and standby management nodes.

**Figure 7-1** Viewing and recording the IP addresses

Hosts

| <input type="checkbox"/> | Host Name | Management IP Addr... | Service IP Address | Running Status |
|--------------------------|-----------|-----------------------|--------------------|----------------|
| <input type="checkbox"/> | 1         |                       |                    | ● Normal       |
| <input type="checkbox"/> | 2         |                       |                    | ● Normal       |
| <input type="checkbox"/> | 3         |                       |                    | ● Normal       |
| <input type="checkbox"/> | ★ 7       |                       |                    | ● Normal       |
| <input type="checkbox"/> | ★ 8       |                       |                    | ● Normal       |
| <input type="checkbox"/> | 9         |                       |                    | ● Normal       |

- If a host name starts with ★, it is the active management node (active OMS node). View and record the value of **Management IP Address** in the row containing the active node.
- If a host name starts with ☆, the host is a standby management node (standby OMS node). View and record the value of **Management IP Address** in the row containing the standby node.

----End

## 7.2.2 Logging In to an MRS Cluster Node

This section describes how to remotely log in to an ECS in an MRS cluster using the remote login (VNC mode) function provided on the ECS management console or a key or password (SSH mode). Remote login (VNC mode) is mainly used for emergency O&M. In other scenarios, it is recommended that you log in to the ECS using SSH.

**NOTE**

To log in to a cluster node using SSH, you need to add an inbound rule to the security group of the cluster. Set **Source** to *IPv4 address of the client/32* or *IPv6 address of the client/128* and set the port number to **22**. For details, see [Adding a Security Group Rule](#).

## Logging In to an ECS Using VNC

- Step 1** Log in to the MRS management console.
- Step 2** Choose **Active Clusters** in the navigation pane, select a running cluster, and click its name to switch to the cluster details page.
- Step 3** On the **Nodes** tab page, click the name of a Master node in the Master node group to log in to the ECS management console.
- Step 4** In the upper right corner, click **Remote Login**.
- Step 5** Enter the username and password for logging in to the Master node as prompted.
  1. If you select **Password** for **Login Mode**, you need to enter **root** in **Username** and the password you set during cluster creation in **Password**.

**Figure 7-2** Selecting password as the login mode

\* Login Mode: Password (selected), Key Pair

Username: root

Keep your password secure. The system cannot retrieve your password.

\* Password: [input field]

\* Confirm Password: [input field]

2. If you select **Key Pair** for **Login Mode** when creating a cluster, perform the following operations to log in to the cluster:
  - a. After the cluster is created, assign an EIP and bind it to the Master node of the cluster. For details, see [Assigning an EIP and Binding It to an ECS](#).
  - b. Remotely log in to the Master node in SSH mode as user **root** using the key file.
  - c. Run the **passwd root** command to set a password for user **root**.
  - d. Go back to the login interface, and enter **root** and the password set in [Step 5.2.c](#) to log in to the node.

----End

## Logging In to an ECS Using a Key Pair (SSH)

### Logging In to the ECS from Local Windows

To log in to the Linux ECS from local Windows, perform the operations described in this section. The following procedure uses PuTTY as an example to log in to the ECS.

**Step 1** Log in to the MRS management console.

**Step 2** Choose **Active Clusters** in the navigation pane, select a running cluster, and click its name to switch to the cluster details page.

**Step 3** On the **Nodes** tab page, click the name of a Master node in the Master node group to log in to the ECS management console.

**Step 4** Click the **EIPs** tab, click **Bind EIP** to bind an EIP to the ECS, and record the EIP. If an EIP has been bound to the ECS, skip this step.

**Step 5** Check whether the private key file has been converted to **.ppk** format.

- If yes, go to **Step 10**.
- If no, go to **Step 6**.

**Step 6** Run PuTTY.

**Step 7** In the **Actions** area, click **Load** and import the private key file you used during ECS creation.

Ensure that the private key file is in the format of **All files (\*.\*)**.

**Step 8** Click **Save private key**.

**Step 9** Save the converted private key, for example, **kp-123.ppk**, to a local directory.

**Step 10** Run PuTTY.

**Step 11** Choose **Connection > Data**. Enter the image username in **Auto-login username**.

 **NOTE**

The image username for cluster nodes is **root**.

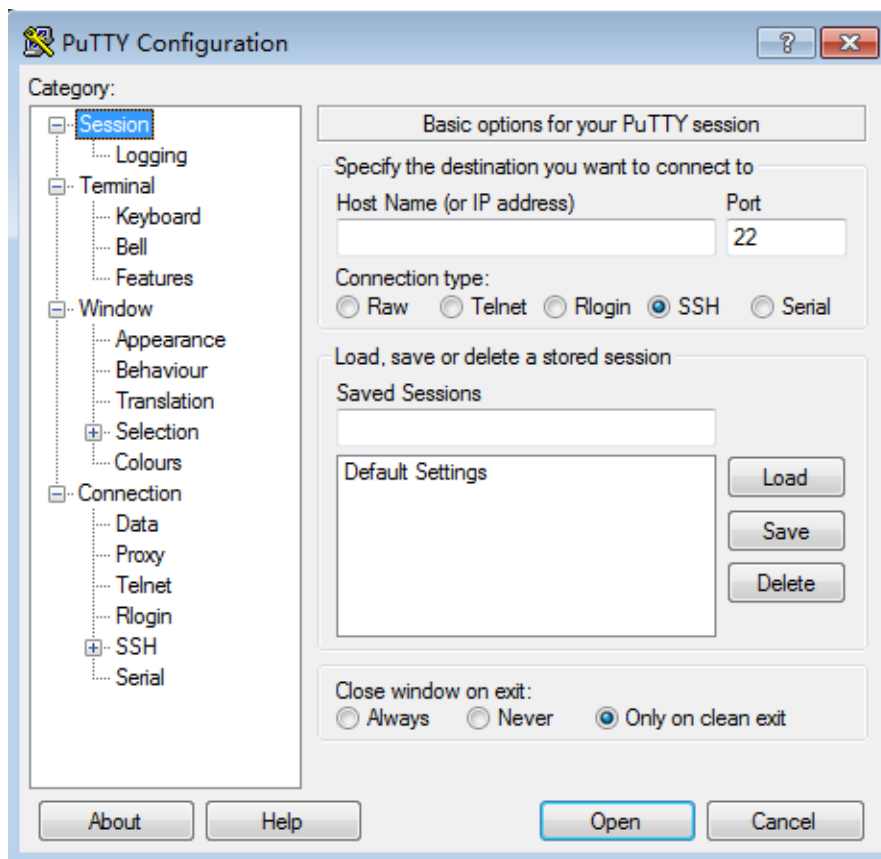
**Step 12** Choose **Connection > SSH > Auth**. In the last configuration item **Private key file for authentication**, click **Browse** and select the private key in .ppk format or the private key saved in **Step 9**.

**Step 13** Click **Session**.

1. **Host Name (or IP address)**: Enter the EIP bound to the ECS.
2. **Port**: Enter **22**.
3. **Connection Type**: Select **SSH**.
4. **Saved Sessions**: Task name, which can be clicked for remote connection when you use PuTTY next time



Figure 7-3 Clicking Session



**Step 14** Click **Open** to log in to the ECS.

If you log in to the ECS for the first time, PuTTY displays a security warning dialog box, asking you whether to accept the ECS security certificate. Click **Yes** to save the certificate to your local registry.

----End

### Logging In to the ECS from Local Linux

To log in to the Linux ECS from local Linux, perform the operations described in this section. The following procedure uses private key file **kp-123.pem** as an example to log in to the ECS. The name of your private key file may differ.

**Step 1** On the Linux CLI, run the following command to change operation permissions:

```
chmod 400 /path/kp-123.pem
```

#### NOTE

In the preceding command, replace *path* with the actual path where the key file is saved.

**Step 2** Run the following command to log in to the ECS:

```
ssh -i /path/kp-123.pem Default username@EIP
```

For example, if the default username is **root** and the EIP is **123.123.123.123**, run the following command:

```
ssh -i /path/kp-123.pem root@123.123.123.123
```

 **NOTE**

- *path* indicates the path where the key file is saved.
- *EIP* indicates the EIP bound to the ECS.
- The image username is **root** for cluster nodes.

----End


## Logging In to an ECS Using a Password (SSH)

### Logging In to the ECS from Local Windows

To log in to the Linux ECS from local Windows, perform the operations described in this section. The following procedure uses PuTTY as an example to log in to the ECS.

- Step 1** Log in to the MRS management console.
- Step 2** Choose **Active Clusters** in the navigation pane, select a running cluster, and click its name to switch to the cluster details page.
- Step 3** On the **Nodes** tab page, click the name of a Master node in the Master node group to log in to the ECS management console.
- Step 4** Click the **EIPs** tab, click **Bind EIP** to bind an EIP to the ECS, and record the EIP. If an EIP has been bound to the ECS, skip this step.
- Step 5** Run PuTTY.
- Step 6** Click **Session**.
  1. **Host Name (or IP address)**: Enter the EIP bound to the ECS.
  2. **Port**: Enter **22**.
  3. **Connection Type**: Select **SSH**.
  4. **Saved Sessions**: Task name, which can be clicked for remote connection when you use PuTTY next time

**Figure 7-4** Clicking **Session**



**Step 7** Click **Window** and select **UTF-8** for **Remote character set:** in **Translation**.

**Step 8** Click **Open** to log in to the ECS.

If you log in to the ECS for the first time, PuTTY displays a security warning dialog box, asking you whether to accept the ECS security certificate. Click **Yes** to save the certificate to your local registry.

**Step 9** After the SSH connection to the ECS is set up, enter the username and password as prompted to log in to the ECS.

**NOTE**

The username is **root** and the password is the one you set during cluster creation.

----End

**Logging In to the ECS from Local Linux**

If the local host runs Linux, perform steps **Step 1** to **Step 4** to bind an EIP to the ECS, and run the following command on the CLI to log in to the ECS: **ssh EIP bound by the ECS**

## 7.3 Viewing MRS Cluster Monitoring Metrics

## 7.3.1 Viewing MRS Cluster Resource Monitoring Metrics

MRS cluster nodes are classified into management nodes, control nodes, and data nodes. The change trends of key host monitoring metrics on each type of node can be calculated and displayed as curve charts in reports based on the customized periods. MRS cluster metrics are monitored periodically. The average historical monitoring interval is about 5 minutes.

You can view the overall resource overview of a cluster on the MRS management console or Manager.

### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Viewing Monitoring Information on the MRS Management Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the **Dashboard** tab, click **Synchronize** next to **IAM User Sync** to synchronize IAM users.

**Step 4** After the synchronization is complete, click the **Monitor** tab to view the monitoring metric report of the cluster.

**Step 5** In time range area, specify a period to view monitoring data.

**Step 6** Customize a monitoring report.

1. Click **Customize** and select monitoring metrics to display.
2. Click **OK** to save the selected monitoring metrics for display.

#### NOTE

Click **Clear All** to deselect all selected metrics in batches.

**Step 7** Export a monitoring report.


1. Select a period.
2. Click **Export**. MRS will generate a report about the selected monitoring metrics in a specified time range. Save the report.

----End

### Viewing Monitoring Information on Manager (MRS 3.x and Later)

**Step 1** Log in to Manager.

**Step 2** Choose **Homepage**.

**Step 3** In the upper right corner of the chart area, click  and choose **Customize** from the displayed menu.

 **NOTE**




Monitoring data of the past 1 hour is displayed at an interval of 5 minutes. After you enter the **Real-time Monitoring** page, you can view that real-time monitoring data is displayed on the right of the monitoring chart at an interval of 5 minutes.


**Step 4** In the navigation pane, select an entity and some monitoring metrics.

**Step 5** Click **OK** to view the metric details.

**Step 6** Export a monitoring report.

- Exporting all monitoring data
  - a. On the **Homepage**, select a time range in the upper right corner of the chart area, for example, **1w**.

Real-time data is displayed by default, which cannot be exported. You can click  to customize a time range.
  - b. In the upper right corner of the chart area, click  and choose **Export** from the displayed menu.
- Exporting monitoring data of a specified monitoring item
  - a. On the **Homepage**, click  in the upper right corner of any monitoring report pane in the chart area of the target cluster.
  - b. Select a time range to obtain monitoring data, for example, **1w**.

Real-time data is displayed by default, which cannot be exported. You can click  to customize a time range.
  - c. Click **Export**.

 **NOTE**

The interval on the horizontal axis of the chart varies depending on the time period you specify. Data monitoring rules are as follows:

- If the disk usage of the partition where GaussDB is deployed exceeds 80%, real-time monitoring data and monitoring data whose interval is 5 minutes will be deleted.
- **Storage resources (HDFS) in Tenant Resources (0 to 300 hours)**: The interval is 1 hour. The cluster must have been installed for at least 1 hour, and monitoring data of a maximum of 3 months is saved.
- **0 to 25 hours**: The interval is 5 minutes. The cluster must have been installed for at least 10 minutes, and monitoring data of a maximum of 15 days is saved.
- **25 to 150 hours**: The interval is 30 minutes. The cluster must have been installed for at least 30 minutes, and monitoring data of a maximum of 3 months is saved.
- **150 to 300 hours**: The interval is 1 hour. The cluster must have been installed for at least 1 hour, and monitoring data of a maximum of 3 months is saved.
- **300 hours to 300 days**: The interval is 1 day. The cluster must have been installed for at least 1 day, and monitoring data of a maximum of 6 months is saved.
- **Over 300 days**: The interval is 7 days. The cluster must have been installed for more than 7 days, and monitoring data of a maximum of 1 year is saved.

----End


## Viewing Monitoring Information on Manager (MRS 2.x and Earlier)

**Step 1** Log in to Manager and click **System**.

**Step 2** In **Period**, you can specify a period to view monitoring data. The options are as follows:

Real-time, last 3 hours, last 6 hours, last 24 hours, last 1 week, last 1 month, last 3 months, last 6 months, and custom time range


**Step 3** Click **View** to view monitoring data in a specified period.

- You can view **Health Status** and **Roles** of each service on the **Service Summary** page of MRS Manager.
- Click  above the curve chart to view details about a metric.

**Step 4** Customize a monitoring report.

Click **Customize** and select monitoring metrics to be displayed on MRS Manager. Click **OK** to save the selected monitoring metrics for display. Click **Clear** to cancel all the selected monitoring metrics in a batch.

MRS Manager supports a maximum of 14 monitoring metrics, but at most 12 customized monitoring metrics can be displayed on the page.

**Step 5** Set an automatic refresh interval or click  for an immediate refresh. If you select **Full Screen**, the **Dashboard** window will be maximized.

The following refresh interval options are supported:

- **Refresh every 60 seconds**
- **Refresh every 120 seconds**
- **Stop refreshing**

**Step 6** Export a monitoring report.

Select a period. Click **Export**. MRS Manager will generate a report about the selected monitoring metrics in a specified time of period. Select a directory to save the report.

To view the curve charts of monitoring metrics in a specified period, click **View**.

 **NOTE**

You can query top, bottom, and average value curves for key service and host monitoring metrics on MRS Manager, which displays resource distribution information. MRS Manager also shows monitoring data for the last hour.

You can also modify the resource distribution on MRS Manager to display both the top and bottom value curves in service and host resource distribution figures.

Resource distribution of some monitoring metrics is not recorded.

- View the resource distribution of service monitoring metrics.
  1. On MRS Manager page, click **Services**.
  2. Select the target service from the service list.
  3. Click **Resource Distribution**.

Select key metrics of the service from **Metric**. MRS Manager displays the resource distribution of the metrics in the last hour.
- View the resource distribution of host monitoring metrics.
  1. On MRS Manager, click **Hosts**.
  2. Click the name of the specified host in the host list.
  3. Click **Resource Distribution**.

Select key metrics of the host from **Metrics**. MRS Manager displays the resource distribution of the metrics in the last hour.
- Configure resource distribution.
  1. On MRS Manager, click **System**.
  2. In **Configuration**, click **Configure Resource Contribution Ranking** under **Monitoring and Alarm**.
  3. Change the number of resources to be displayed. The sum of the maximum value and minimum value of resource distribution cannot be greater than 5.

Set **Number of Top Resources** to the number of top values.

Set **Number of Bottom Resources** to the number of bottom values.
  4. Click **OK** to save the configurations.

The message "Number of top and bottom resources saved successfully" is displayed in the upper right corner of the page.

----End

## 7.3.2 Viewing MRS Cluster Component Monitoring Metrics

You can manage the following status and metrics of all components (including role instances) on the MRS console: Status information includes operation, health, configuration, and role instance status. Indicator information includes key monitoring indicators for each component.

### Prerequisites



- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

## Viewing the Monitoring Information on the MRS Management Console

- Step 1** Log in to the MRS console.
- Step 2** Choose **Active Clusters** and click a cluster name to go to the cluster details page.
- Step 3** On the **Dashboard** tab, click **Synchronize** next to **IAM User Sync** to synchronize IAM users.
- Step 4** On the MRS cluster details page, click **Components**.
- Step 5** View component monitoring information.
1. Click a specified service in the list to view its status and metric information.
  2. Select and view component-level monitoring metrics.
    - a. In the **Charts** area, click **Customize** to customize service monitoring metrics.
    - b. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.
- Step 6** View role instance monitoring information.
1. In the component list, click a service name.
  2. Click **Instance** to view the status of each role instance in the component.  
You can filter all instances of the same role in the upper right corner of the list. You can set search criteria in the role search area by clicking **Advanced Search**, and click **Search** to view specified role information. You can click **Reset** to reset the search criteria. Fuzzy search is supported.
  3. Click the target role instance to view its status and metric information.
  4. Customize and view monitoring graphs.
    - a. In the **Charts** area, click **Customize** to customize service monitoring metrics.
    - b. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.


----End

## Viewing Component Monitoring Information on Manager

- Step 1** Log in to Manager.
- Step 2** Go to the **Service Management** page.
- For MRS 3.x and later, choose **Cluster > Services**.
  - For MRS 2.x and earlier, click **Services**.
- Step 3** View component monitoring information.
1. Click a specified service in the list to view its status and metrics.
  2. Customize and export monitoring charts.
    - a. In the **Chart** tab, click  and **Customize** to customize service monitoring metrics.
    - b. Set a time range and click  and **Export** to export the monitoring data.



**Step 4** View role instance monitoring information.

1. Click the service name.
2. Click **Instances** to view the role status.
3. Click the target role instance to view its status and metric information.
4. Customize and export monitoring charts.
  - a. In the **Chart** tab, click  and **Customize** to customize metric chart.
  - b. Select a time range. The monitoring data within the time range is displayed.



----End

## Component Resource Monitoring Summary


 **NOTE**

This function is supported only in MRS 3.x and later versions.

Log in to FusionInsight Manager and choose **Cluster > Services > Target service**. Click **Resource**. The resource monitoring page is displayed.

Some services in the cluster provide service-level resource monitoring metrics. By default, the monitoring data of the latest 12 hours is displayed. You can click  to set a time range. You can click  to export the corresponding report information. If a monitoring item has no data, the report cannot be exported. The following table lists the services and monitoring items that support resource monitoring.

**Table 7-2** Service resource monitoring

| Service | Monitoring Metric             | Description                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HDFS    | Resource Usage (by Tenant)    | <ul style="list-style-type: none"> <li>• Collects statistics on HDFS resource usage by tenant.</li> <li>• Views the metrics <b>Capacity</b> or <b>Number of File Objects</b>.</li> </ul>                                                                                                                                                                                                                                    |
|         | Resource Usage (by User)      | <ul style="list-style-type: none"> <li>• Collects statistics on HDFS resource usage by user.</li> <li>• Views the metrics <b>Used Capacity</b> or <b>Number of File Objects</b>.</li> </ul>                                                                                                                                                                                                                                 |
|         | Resource Usage (by Directory) | <ul style="list-style-type: none"> <li>• Collects statistics on HDFS resource usage by directory.</li> <li>• Views the metrics <b>Used Capacity</b> or <b>Number of File Objects</b>.</li> <li>• You can click  to configure space monitoring. Alternatively, you can specify an HDFS file system directory for monitoring.</li> </ul> |

| Service    | Monitoring Metric                       | Description                                                                                                                                                                                                  |
|------------|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|            | Resource Usage (by Replica)             | <ul style="list-style-type: none"> <li>Collects statistics on HDFS resource usages by replica count.</li> <li>Views the metrics <b>Used Capacity</b> or <b>File Count</b>.</li> </ul>                        |
|            | Resource Usage (by File Size)           | <ul style="list-style-type: none"> <li>Collects statistics on HDFS resource usages by file size.</li> <li>Views the metrics <b>Used Capacity</b> or <b>File Count</b>.</li> </ul>                            |
|            | Recycle Bin (by User)                   | <ul style="list-style-type: none"> <li>Collects statistics on the usage of the HDFS recycle bin by user.</li> <li>Views the metrics <b>Recycle Bin Capacity</b> or <b>Number of File Objects</b>.</li> </ul> |
|            | Operation Count                         | <ul style="list-style-type: none"> <li>Collects the number of operations in HDFS.</li> </ul>                                                                                                                 |
|            | Automatic Balancer                      | Collects statistics on the execution speed of HDFS automatic balancer and the total capacity of the current balancer migration.                                                                              |
|            | NameNode RPC Open Connections (by User) | Displays the number of connections of each user in the Client RPC requests connected to NameNodes.                                                                                                           |
|            | Slow DataNodes                          | Displays DataNode that transmits or processes data slowly in the cluster.                                                                                                                                    |
|            | Slow Disks                              | Displays the disk that processes data slowly on the DataNode in the cluster.                                                                                                                                 |
| HBase      | Operation Requests in Tables            | Displays the number of PUT, DELETE, GET, SCAN, INCREMENT, and APPEND operation requests in all tables on all RegionServers.                                                                                  |
|            | Operation Requests on RegionServers     | Displays the number of PUT, DELETE, GET, SCAN, INCREMENT, and APPEND operation requests and number of all operation requests in RegionServer.                                                                |
|            | Operation Requests for Service          | Displays the number of PUT, DELETE, GET, SCAN, INCREMENT, and APPEND operation requests in all regions on RegionServers.                                                                                     |
|            | HFiles on RegionServers                 | Displays the number of HFiles in all RegionServers.                                                                                                                                                          |
| HetuEngine | Coordinator Resource Usage              | Displays the coordinator resource usage in the selected queue.                                                                                                                                               |

| Service | Monitoring Metric                           | Description                                                                                                                                                                                                                                                                  |
|---------|---------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|         | Coordinator Resource Usage Ratio            | Displays the coordinator resource usage in the selected queue.                                                                                                                                                                                                               |
|         | Worker Resource Usage                       | Displays the worker resource usage in the selected queue.                                                                                                                                                                                                                    |
|         | Worker Resource Usage Ratio                 | Displays the worker resource usage in the selected queue.                                                                                                                                                                                                                    |
|         | Number of Coordinators and Workers          | Displays the number of coordinators and workers in the selected queue.                                                                                                                                                                                                       |
| Hive    | HiveServer2-Background-Pool Threads (by IP) | Displays the number of HiveServer2-Background-Pool threads of top users. These threads are measured and displayed in a measurement period.                                                                                                                                   |
|         | HiveServer2-Handler-Pool Threads (by IP)    | Displays the number of HiveServer2-Handler-Pools of top users collected and displayed in a period.                                                                                                                                                                           |
|         | Used MetaStore Number (by IP)               | Collects statistics on and displays the MetaStore usage of top users in a period.                                                                                                                                                                                            |
|         | Number of Hive jobs                         | Displays the number of user-related jobs collected by Hive in a period.                                                                                                                                                                                                      |
|         | Number of Files Accessed in the Split Phase | Displays the number of files accessed by the underlying file storage system (HDFS by default) in the Split phase in a period.                                                                                                                                                |
|         | Hive Basic Operation Time                   | Collects time for creating a directory (mkdirTime), creating a file (touchTime), writing a file (writeFileTime), renaming a file (renameTime), moving a file (moveTime), deleting a file (deleteFileTime), and deleting a directory (deleteCatalogTime) in a period of time. |
|         | Table Partitions                            | Displays the number of partitions in all Hive tables, which is displayed in the following format: <i>database # table name, number of table partitions</i> .                                                                                                                 |
|         | HQL Map Count                               | Collects statistics on HQL statements executed in a period and the number of Map statements invoked during the execution. The displayed information includes users, HQL statements, and the number of Map statements.                                                        |

| Service           | Monitoring Metric                              | Description                                                                                                                                                                                                                                    |
|-------------------|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                   | HQL Access Statistics                          | Displays the number of HQL access times in a period.                                                                                                                                                                                           |
| Kafka             | Kafka Disk Usage Distribution                  | Displays the disk usage distribution statistics of the Kafka cluster.                                                                                                                                                                          |
| Spark/<br>Spark2x | HQL Access Statistics                          | Collects HQL access statistics in a period, including the username, HQL statement, and HQL statement execution times.                                                                                                                          |
| Yarn              | Used resources (by Task)                       | <ul style="list-style-type: none"> <li>Displays the number of CPU cores and memory used by a task.</li> <li>Views the metrics <b>By memory</b> or <b>By CPU</b>.</li> </ul>                                                                    |
|                   | Resource Usage (by Tenant)                     | <ul style="list-style-type: none"> <li>Displays the number of CPU cores and memory used by a tenant.</li> <li>Views the metrics <b>By memory</b> or <b>By CPU</b>.</li> </ul>                                                                  |
|                   | Resource usage ratio (by Tenant)               | <ul style="list-style-type: none"> <li>Displays the ratio of the number of CPU cores to the memory used by a tenant.</li> <li>Views the metrics <b>By memory</b> or <b>By CPU</b>.</li> </ul>                                                  |
|                   | Task Duration Ranking                          | Displays Yarn tasks sorted by time consumption.                                                                                                                                                                                                |
|                   | ResourceManager RPC Open Connections (by User) | Displays the number of client RPC connections to ResourceManager by user.                                                                                                                                                                      |
|                   | HBase Operation Count                          | Collects statistics on the number and proportion of operations corresponding to each Yarn operation type.                                                                                                                                      |
|                   | Ranking of Tasks in a Queue by Resource Usage  | <ul style="list-style-type: none"> <li>Displays the resources consumed by the tasks running in a queue after the queue (tenant) is selected on the GUI.</li> <li>Views the metrics <b>By memory</b> or <b>By CPU</b>.</li> </ul>               |
|                   | Ranking of Users in a Queue by Resource Usage  | <ul style="list-style-type: none"> <li>Displays the resources consumed by the users who are running tasks in the queue after a queue (tenant) is selected on the GUI.</li> <li>Views the metrics <b>By memory</b> or <b>By CPU</b>.</li> </ul> |
| ZooKeeper         | Used Resources (By Second-Level Znode)         | <ul style="list-style-type: none"> <li>Displays the ZooKeeper level-2 znode resource status.</li> <li>Views the metrics <b>By Znode quantity</b> or <b>By capacity</b>.</li> </ul>                                                             |

| Service | Monitoring Metric                            | Description                                               |
|---------|----------------------------------------------|-----------------------------------------------------------|
|         | Number of Connections (by Client IP Address) | Displays the ZooKeeper client connection resource status. |

### 7.3.3 Viewing MRS Node Resource Monitoring Metrics

You can view the status and metrics of each node in the MRS cluster in real time to see the current resource usage.

#### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

#### Viewing Host Monitoring Charts

##### Operations on the MRS management console:

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.
- Step 3** Click **Nodes** and expand the node group information to view the status of all hosts.
- Step 4** The host list of a group contains the **Node Name/Resource ID, IP, Status, Specifications, Disks, and AZ**.
- Step 5** Click the target node in the list to view its status and metric information.
- Step 6** Click the **Monitoring** tab to view the monitoring charts of the current node.

----End

##### Operations on Manager:

- MRS 3.x and later:
  - a. Log in to FusionInsight Manager.
  - b. Click **Hosts** to view the host list.
  - c. In the host list, click the specified host name to view the host overview information.

The host details page contains the basic information area, disk status area, instance list area, and monitoring charts.
  - d. Click the **Chart** tab to view the full monitoring chart information about the host.

The **Chart** page displays all monitoring charts of the host.


- MRS 2.x and earlier:
  - a. Log in to MRS Manager.
  - b. Click **Hosts** to view the status of all hosts.
  - c. Click the target host in the host list to view its status and metric information.
  - d. Customize and export monitoring charts.
    - i. In the **Charts** area, click **Customize** to customize service monitoring metrics.
    - ii. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.
    - iii. Click **Export** to export the displayed metrics.

## Viewing Host Resource Overview

This function is only available in MRS 3.x or later.

**Step 1** Log in to FusionInsight Manager and choose **Hosts > Resource Overview**.

**Step 2** View the distribution information.

In the **Distribution** tab, you can view the resource distribution of the cluster. By default, the monitoring data of the last hour is displayed. You can click  to set a time range.

**Figure 7-5** Distribution tab



- Click **Select Metric** to specify the metrics you want to check. After you select a metric, the host distribution in each range of the metric is displayed.
- When you place the cursor on a color column, the number of hosts in the range is displayed. You can click a color column to view the list of hosts in the metric range.
  - You can click a host name in the **Host Name** column to access the host details page.
  - Click **View Trends** of a host in the **Operation** column of the list. The metric trends will be displayed. In the current cluster, if you have selected **Host CPU-Memory-Disk Usage**, **View Trends** is unavailable.


- You can click **Export Data** to export the maximum, minimum, and average values of the current metric of all nodes in the cluster within the time range you have specified.

**Table 7-3** Metrics

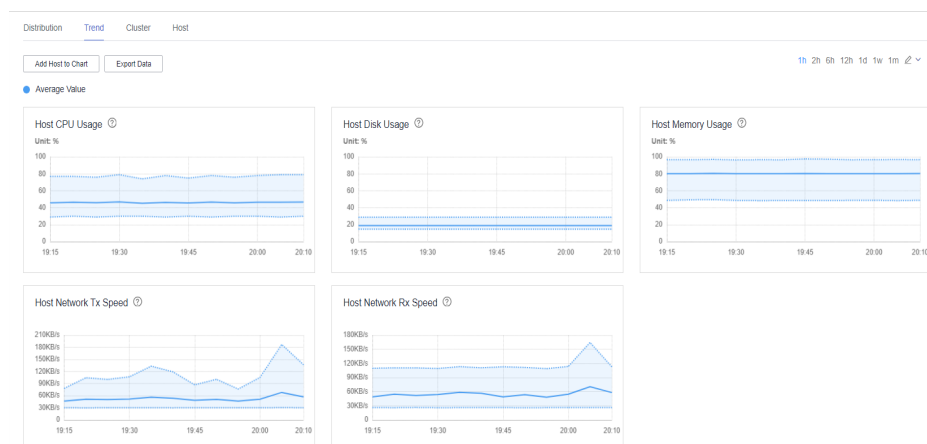
| Category        | Metric                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Process         | <ul style="list-style-type: none"> <li>• Number of Running Processes</li> <li>• Total Number of Processes</li> <li>• Total Number of omm Processes</li> <li>• Total Number of Processes in D and Z States</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                  |
| Network Status  | <ul style="list-style-type: none"> <li>• Host Network Packet Collisions</li> <li>• Number of LAST_ACK States</li> <li>• Number of CLOSING States</li> <li>• Number of LISTENING States</li> <li>• Number of CLOSED States</li> <li>• Number of ESTABLISHED States</li> <li>• Number of SYN_RECV States</li> <li>• Number of TIME_WAITING States</li> <li>• Number of FIN_WAIT2 States</li> <li>• Number of FIN_WAIT1 States</li> <li>• Number of CLOSE_WAIT States</li> <li>• DNS Name Resolution Duration</li> <li>• TCP Ephemeral Port Usage</li> <li>• Host Network Packet Frame Errors</li> </ul> |
| Network Reading | <ul style="list-style-type: none"> <li>• Host Network Read Packets</li> <li>• Host Network Read Dropped Packets</li> <li>• Host Network Read Error Packets</li> <li>• Host Network Receiving Rate</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                          |
| Disk            | <ul style="list-style-type: none"> <li>• Host Disk Write Rate</li> <li>• Host Used Disk</li> <li>• Host Free Disk</li> <li>• Host Disk Read Rate</li> <li>• Host Disk Usage</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                |
| Memory          | <ul style="list-style-type: none"> <li>• Free Memory</li> <li>• Cache Memory Size</li> <li>• Total Kernel Cache Memory Size</li> <li>• Shared Memory Size</li> <li>• Host Memory Usage</li> <li>• Used Memory</li> </ul>                                                                                                                                                                                                                                                                                                                                                                              |


| Category        | Metric                                                                                                                                                                                                                                                                                 |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Network Writing | <ul style="list-style-type: none"> <li>● Host Network Write Packets</li> <li>● Host Network Write Error Packets</li> <li>● Host Network Sending Rate</li> <li>● Host Network Write Dropped Packets</li> </ul>                                                                          |
| CPU             | <ul style="list-style-type: none"> <li>● CPU Usage of Processes Whose Priorities Have Been Changed</li> <li>● CPU Usage of User Space Processes</li> <li>● CPU Usage of Kernel Space Processes</li> <li>● Host CPU Usage</li> <li>● CPU Total Time</li> <li>● CPU Idle Time</li> </ul> |
| Host Status     | <ul style="list-style-type: none"> <li>● Host File Handle Usage</li> <li>● Average OS Load in 1 Minute</li> <li>● Average OS Load in 5 Minutes</li> <li>● Average OS Load in 15 Minutes</li> <li>● Host PID Usage</li> </ul>                                                           |

**Step 3** Click the **Trend** tab to view host monitoring trends.

The resource monitoring trend page of the cluster is displayed. By default, the monitoring data of the past one hour (**1h**) is displayed. You can click  to set a time range.

**Figure 7-6** Trend tab




- You can click **Add Host to Chart** to add trend lines of up to 12 hosts to the trend charts.
- Click  and **Customize** to select the metrics to be displayed on the page.



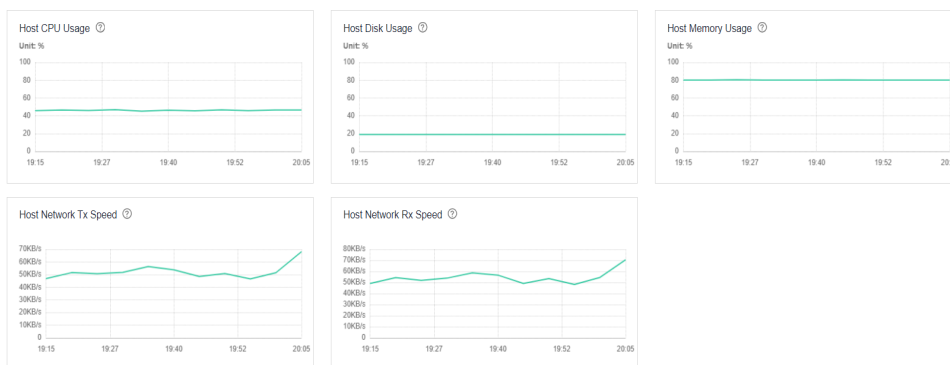
- You can click **Export Data** to export the maximum, minimum, and average values of all nodes in the cluster for all selected metrics within the time range you have specified.


**Step 4** Click the **Cluster** tab to view the cluster information.

You can view the resource monitoring page of each cluster on FusionInsight Manager at the same time.

By default, the monitoring data of the past one hour (1h) is displayed. You can click  to set a time range.

**Figure 7-7** Cluster tab




- Click  and **Customize** to select the metrics to be displayed on the page.
- You can click **Export Data** to export the metric values of each cluster within the time range you have specified.

**Step 5** Click the **Host** tab to view the host information.

You can view the host resource status, including basic configuration (CPU/ memory) and disk configuration.

Click **Export Data** to export the configuration list of all hosts in the cluster.

**Figure 7-8** Host tab



- Basic Configurations (CPU/Memory)

You can hover your cursor over the pie chart to view the number of hosts of each hardware configuration in the cluster. The information is displayed in the format of *Number of cores (CPU architecture) Memory size*.

You can click a slice on the pie chart to view the list of hosts.

- **Disk Configurations**

The horizontal axis indicates the total disk capacity (including the OS disk) of a node, and the vertical axis indicates the number of logical disks (including the OS disk).

You can hover your cursor over a dot to view information about disks of the current configuration, including the quantity of disks, total capacity, and number of hosts.

You can click a dot on the chart to view the list of hosts.

----End

## 7.3.4 Dumping MRS Cluster Monitoring Data

### Scenarios

The monitoring data reporting function writes the monitoring data collected in the system into a text file and uploads the file to a specified server in FTP or SFTP mode. You can configure interconnection parameters on Manager to save monitoring data to a specified FTP server. In this way, MRS clusters can interconnect with third-party systems.

The FTP protocol does not encrypt data, which poses potential security risks. Therefore, the SFTP protocol is recommended. The ECS corresponding to the dump server must be in the same VPC as the Master nodes of the MRS cluster, and the Master nodes can access the IP address and specified port of the dump server. The FTP service on the dump server is running properly.

Manager supports the collection of all the monitoring data in the managed clusters. The collection period is 30 seconds, 60 seconds, or 300 seconds. The monitoring data is stored in different monitoring files on the FTP server by collection period.

**Table 7-4** Monitoring data files

| Monitoring Data   | Description                                                                                                                                                                                                                                                   |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Collection period | <ul style="list-style-type: none"><li>• 30s: real-time metrics that are collected every 30s by default</li><li>• 60s: real-time metrics that are collected every 60s by default</li><li>• 300s: all metrics that are not collected every 30s or 60s</li></ul> |

| Monitoring Data                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| File name format                       | <ul style="list-style-type: none"> <li>MRS 3.x and later:<br/><b>metric_Monitoring data collection period_File creation time.log</b></li> <li>Examples: <b>metric_60_20160908085915.log</b> and <b>metric_300_20160908085613.log</b></li> <li>Versions earlier than MRS 3.x:<br/><i>Cluster name_metric_Monitoring data collection period_File creation time.log</i></li> </ul>                                                                                                                                             |
| Content format (MRS 3.x as an example) | <ul style="list-style-type: none"> <li>Cluster ID Cluster name Display name Service name  <b>Metric ID</b> Collection time Collection host@m@Sub-metric Unit Metric value</li> <li>Example:<br/>The actual files do not contain the parameter names in the format description.</li> </ul> <pre>1 xx1 Host Host 10000413 2019/06/18 10:05:00 10-66-254-146 KB/s 309.910 1 xx1 Host Host 10000413 2019/06/18 10:05:00 10-66-254-152 KB/s 72.870 2 xx2 Host Host 10000413 2019/06/18 10:05:00 10-66-254-163 KB/s 100.650</pre> |

With the **Metric ID** obtained from the reported file, the third-party can obtain metric details by querying the metric set file in the *FusionInsight installation path/om-server/om/etc/om/all-shown-metric-zh\_CN* on the active/standby OMS node. The file contains the detailed information about all metrics. The content is as follows (taking MRS 3.x as an example):

```
Real-Time Metric ID,5-Minute Metric ID,Metric Name,Metric Collection Period (s),Collected by
Default,Service Belonged To,Role Belonged To
00101,10000101,JobHistoryServer non-heap memory usage,30,false,Mapreduce,JobHistoryServer
00102,10000102,JobHistoryServer non-heap memory allocation volume,30,false,Mapreduce,JobHistoryServer
00103,10000103,JobHistoryServer heap memory usage,30,false,Mapreduce,JobHistoryServer
00104,10000104,JobHistoryServer heap memory allocation volume,30,false,Mapreduce,JobHistoryServer
00105,10000105,Number of blocked threads,30,false,Mapreduce,JobHistoryServer
00106,10000106,Number of running threads,30,false,Mapreduce,JobHistoryServer
00107,10000107,GC time,30,false,Mapreduce,JobHistoryServer
00110,10000110,JobHistoryServer CPU usage,30,false,Mapreduce,JobHistoryServer
...
```

- Metric ID meaning:
  - For metrics whose collection period is 30s/60s, you can find the corresponding metric description by referring to the first column, that is, **Real-Time Metric ID**.
  - For metrics whose collection period is 300s, you can find the corresponding metric description by referring to the second column, that is, **5-Minute Metric ID**.
- Field description:
  - Real-Time Metric ID**: indicates the ID of the metric whose collection period is 30s or 60s.
  - 5-Minute Metric ID**: indicates the ID of a 5-minute (300s) metric.
  - Metric Collection Period (s)**: indicates the collection period of real-time metrics. The value can be **30** or **60**.

**Service Belonged To:** indicates the name of the service to which a metric belongs, for example, HDFS and HBase.


**Role Belonged To:** indicates the name of the role to which a metric belongs, for example, JobServer and RegionServer.

## Dumping MRS Cluster Monitoring Data (MRS 3.x and Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **System > Interconnection > Upload Performance Data**.

**Step 3** Toggle on **Upload Performance Data**.

The performance data upload service is disabled by default.  indicates that the service is enabled.

**Step 4** Set the upload parameters according to [Table 7-5](#).

**Table 7-5** Upload parameters

| Parameter               | Description                                                                                                                                                                                                                               |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FTP IP Address Mode     | Specifies the server IP address mode. This parameter is mandatory. The value can be <b>IPV4</b> or <b>IPV6</b> .                                                                                                                          |
| FTP IP Address          | Specifies the IP address of the FTP server for storing monitoring files after the monitoring metric data is interconnected. This parameter is mandatory.                                                                                  |
| FTP Port                | Specifies the port for connecting to the FTP server. This parameter is mandatory.                                                                                                                                                         |
| FTP Username            | Specifies the username for logging in to the FTP server. This parameter is mandatory.                                                                                                                                                     |
| FTP Password            | Specifies the password for logging in to the FTP server. This parameter is mandatory.                                                                                                                                                     |
| Save Path               | Specifies the path for storing monitoring files on the FTP server. This parameter is mandatory.                                                                                                                                           |
| Dump Interval (second)  | Specifies the interval at which monitoring files are periodically stored on the FTP server, in seconds. This parameter is mandatory. The system will periodically upload files to the corresponding FTP server at the specified interval. |
| Dump Mode               | Specifies the protocol used for sending monitoring files. This parameter is mandatory. The value can be <b>SFTP</b> or <b>FTP</b> . You are advised to use the SFTP mode based on SSH v2. Otherwise, security risks may be incurred.      |
| SFTP Service Public Key | Specifies the public key of the FTP server. This parameter is optional. It is valid only when <b>Dump Mode</b> is set to <b>SFTP</b> .                                                                                                    |

**Step 5** Click **OK**.

 **NOTE**

If the dump mode is SFTP and the public key of the SFTP service is empty, the system displays a security risk warning. You need to evaluate the security risk and then save the configuration.

----End

## Dumping MRS Cluster Monitoring Data (MRS 2.x and Earlier)

- Step 1** On MRS Manager, click **System**.
- Step 2** In **Configuration**, click **Configure Monitoring Metric Dump** under **Monitoring and Alarm**.
- Step 3** [Table 7-6](#) describes the parameters for dumping.

**Table 7-6** Parameters

| Parameter               | Description                                                                                                                                                                                    | Mandatory |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Dump Monitoring Metric  | Whether to enable monitoring metric data interconnection.                                                                                                                                      | Yes       |
| FTP IP Address          | FTP server for storing monitoring files after monitoring metric data is interconnected.                                                                                                        | Yes       |
| FTP Port                | Port for connecting to the FTP server.                                                                                                                                                         | Yes       |
| FTP Username            | Username for logging in to the FTP server.                                                                                                                                                     | Yes       |
| FTP Password            | Password for logging in to the FTP server.                                                                                                                                                     | Yes       |
| Save Path               | Path for storing monitoring files on the FTP server.                                                                                                                                           | Yes       |
| Dump Interval (second)  | Interval at which monitoring files are periodically stored on the FTP server, in seconds. The system will periodically upload files to the corresponding FTP server at the specified interval. | Yes       |
| Dump Mode               | Protocol used for sending monitoring files. The options are <b>FTP</b> and <b>SFTP</b> .                                                                                                       | Yes       |
| SFTP Service Public Key | Public key of the FTP server. This parameter is available only when <b>Dump Mode</b> is set to <b>SFTP</b> . You are advised to configure a public key. Otherwise, security risks may arise.   | No        |

- Step 4** Click **OK**.

----End

## 7.4 Checking MRS Cluster Health

## 7.4.1 Performing a Health Check for an MRS Cluster

The cluster health check includes three key items: object health, related alarms, and custom metrics. The check result may not match the health status shown on the UI. A health check includes MRS Manager, service-level, and host-level health checks:

- MRS Manager health checks focus on whether the unified management platform can provide management functions properly.
- Service-level health checks focus on whether components can provide services properly.
- Host-level health checks focus on whether host indicators are normal.

To ensure that cluster parameters, configurations, and monitoring are correct and that the cluster can run stably for a long time, you can perform a health check during routine maintenance.

### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Health Check on the MRS Management Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** Start the health check.

#### NOTE

Health check operations on the MRS management console are valid for clusters of **MRS 1.9.2** only.

- Perform the health check for all services.

Choose **Management Operations > Start Cluster Health Check**.

- Perform the health check for a service.

Click **Components**. In the service list, click the target service name and choose **More > Start Service Health Check**.

- Perform the health check for a host.

Click **Nodes**, expand the node group information, select the check box of the target hosts, and choose **Node > Start Host Health Check**.

----End

### Performing a Health Check on Manager

For MRS 3.x and later versions:

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Health Check**.

By default, all saved health check reports are displayed in a list. For details, see [Table 7-7](#).

**Step 3** Start health checks.

- Start a health check.  
Click **Start Check**. In the displayed dialog box, click **OK** to start a health check.
- Set periodic automatic health check.  
Click **Configuration**, select **Enable**, set the check period to **Daily**, **Weekly**, or **Monthly** based on your needs, and click **OK** to save the configuration.

----End

**For MRS 2.x or earlier:**

**Step 1** Log in to MRS Manager.

**Step 2** Start health checks.

- Start a health check.  
Click **Services** and choose **More > Start Service Health Check**.
- Set periodic automatic health check.
  - a. Click **System**, click **Check Health Status** under **Maintenance**, and click **Configure Health Check**.
  - b. Set the **Max. Number of Health Check Reports**. The value must be an integer ranging from 1 to 100.
  - c. Click **Periodic Health Check** to enable this function. Set the check period to **Daily**, **Weekly**, or **Monthly** based on O&M requirements, and click **OK** to save the configuration.

----End

## 7.4.2 Performing Health Checks on MRS Cluster Nodes

If the running status of a host is not **Normal**, you can perform health checks on the host to check whether some basic functions are abnormal. During routine O&M, you can perform host health checks to ensure that the configuration parameters and monitoring of each role instance on the host are normal and can run stably for a long time.

### Performing Health Checks on Cluster Nodes (MRS 3.x and Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Click **Hosts**.

**Step 3** Select the check box of the target host.

**Step 4** Select **Health Check** from the **More** drop-down list to start the health check.

To export the result of the health check, click **Export Report** in the upper left corner. If any problem is detected, click **Help**.

----End

## Performing Health Checks on Cluster Nodes (MRS 2.x and Earlier)

**Step 1** Log in to MRS Manager.

**Step 2** Click **Hosts**.

**Step 3** Select the check box of the host for which you want to check the health status.

**Step 4** Choose **More > Start Host Health Check** to start the health check for the host.

----End

## 7.4.3 Viewing and Exporting a Health Check Report

You can view the health check result on MRS and export it for further analysis.

### Viewing the Health Check Report on the Management Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** On the cluster details page, choose **Management Operations > View Cluster Health Check Report**.

#### NOTE

Health check operations on the MRS management console are valid for clusters of **MRS 1.9.2** only.

**Step 4** Click **Export Report** on the health check report pane to export the report and view detailed information about check items.

----End

### Viewing the Health Check Report on Manager (MRS 3.x and Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Health Check**.

- View the health check report.

All saved check reports are listed by default. You can filter the check object and result status in the upper right corner of the list. If **Check Type** is **Cluster**, **View Help** is displayed in the **Check Object** drop-down list. During a health check, the system determines whether check objects are healthy based on their historical monitoring metric data.



**Table 7-7** Parameters for a health check report

| Item         | Description                                                                                                                                                                      |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Check Object | Object to be checked. You can expand the list to view its details.                                                                                                               |
| Status       | Check result status. Value options are <b>No problems found, Problems found, and Checking.</b>                                                                                   |
| Check Type   | Entity on which the check is to be performed. Value options are <b>System, Cluster, Host, Service, and OMS.</b> If you select <b>Cluster</b> , all items are checked by default. |
| Start Mode   | Whether the health check is automatically or manually performed                                                                                                                  |
| Start Time   | Start time of the check                                                                                                                                                          |
| End Time     | End time of the check                                                                                                                                                            |
| Operation    | Operations you can perform. Value options are <b>Export Report</b> and <b>View Help.</b>                                                                                         |

- Export the health check report.  
Locate the row containing the target health check report and click **Export Report** in the **Operation** to download the report.

----End

## Viewing the Health Check Report on Manager (MRS 2.x and Earlier)

**Step 1** Log in to MRS Manager.

**Step 2** View or export the health check report.

- View the health check report.
  - a. Click **Services**.
  - b. Choose **More > View Cluster Health Check Report** to view the health check report of a cluster. Click **Export Report** on the health check report pane to export the report and view detailed information about check items.
- Download the health check report.
  - a. Choose **System > Maintenance > Check Health Status**.
  - b. Locate the row that contains the target health check report and click **Download** to download the report file.
- Configuring the Number of Health Check Reports to Be Reserved  
Health check reports of MRS clusters, services, and hosts may vary with the time and scenario. You can modify the number of health check reports to be reserved on MRS Manager for later comparison. This setting is valid for health check reports of clusters, services, and hosts. Report files are saved in **\$BIGDATA\_DATA\_HOME/Manager/healthcheck** on the active management node by default and are automatically synchronized to the standby management node.

- a. Choose **System > Maintenance > Check Health Status > Configure Health Check**.
- b. Set **Max. Number of Health Check Reports** to the number of health check reports to be reserved. The value ranges from 1 to 100. The default value is 50.
- c. Click **OK** to save the settings.

----End

## 7.5 Adjusting the Capacity of an MRS Cluster

### 7.5.1 Scaling Out an MRS Cluster

The storage and computing capabilities of MRS can be improved by simply adding Core nodes or Task nodes instead of modifying system architecture, reducing O&M costs. Core nodes can process and store data. You can add Core nodes to expand the node quantities and handle peak loads. Task nodes are used to process data instead of storing persistent data.

#### NOTE

- Only Master, Core, and Task nodes can be added.
- When you log in to a node added for scale-out as user **root**, the password set during cluster creation is required.

### Constraints

- When you expand a node group where HBase is installed:  
If automatic DNS registration is not enabled for a node in the cluster, do not start HBase when you expand the node group. Then, update the HBase client configuration by referring to [Updating the MRS Cluster Client After the Server Configuration Expires](#) and start the HBase instances on the node to be expanded.  
Automatic DNS registration is enabled by default in the following versions: MRS 1.9.3, MRS 3.1.0, MRS 3.1.2-LTS, MRS 3.1.5, and MRS 3.2.0-LTS  
You can check whether this DNS function is supported by checking whether the **features** field in the response body contains **register\_dns\_server**. For details, see [Querying the Metadata of a Cluster Version](#).
- After a scale-out, the clients installed on nodes in the cluster do not need to be updated. For details about how to update the client installed on nodes outside the cluster, see [Updating the MRS Cluster Client After the Server Configuration Expires](#).
- If you need to balance HDFS data after scale-out, see [Balancing DataNode Capacity](#). For details about how to balance Kafka data, see [Kafka Balancing Tool Instructions](#).

### Scaling Out a Cluster Billed in Pay-per-Use Mode

**Step 1** Log in to the MRS console.

- Step 2** Choose **Active Clusters**, select a running cluster, and click its name to switch to the cluster details page.
- Step 3** Click the **Nodes** tab. In the **Operation** column of the node group, click **Scale Out**. The **Scale Out** page is displayed.
- The scale-out operation can only be performed on the running clusters.
- Step 4** Set the type of **System Disk** and **Data Disk**, **Scale-Out Nodes**, **Enable Components** and **Run Bootstrap Action**, and click **OK**. The **Enable Components** and **Run Bootstrap Action** parameters may not be supported by clusters of some versions. The operations are subject to the UI.

×

### Scale Out

|                     |                                            |                                  |                                |
|---------------------|--------------------------------------------|----------------------------------|--------------------------------|
| Node Type           | <input type="text" value="Analysis Core"/> |                                  |                                |
| Node Specifications | 8 vCPUs 32 GB   Sit3.2xlarge.4             |                                  |                                |
| System Disk         | <input type="text" value="High I/O"/>      | <input type="text" value="480"/> | <input type="text" value="+"/> |
| Data Disk           | <input type="text" value="High I/O"/>      | <input type="text" value="600"/> | <input type="text" value="+"/> |
| Disks               | <input type="text" value="1"/>             |                                  |                                |
| Current Nodes       | 3                                          |                                  |                                |
| Scale-Out Nodes     | <input type="text" value="0"/>             |                                  |                                |

Insufficient node quota. [Increase quota](#)

#### NOTE

- If the Task node group does not exist in the cluster, configure the Task node by referring to [Related Tasks](#).
- If a bootstrap action is added during cluster creation, the **Run Bootstrap Action** parameter is valid. If this function is enabled, the bootstrap actions added during cluster creation will be run on all the scaled out nodes.
- If the **New Specifications** parameter is available, the specifications that are the same as those of the original nodes have been sold out or discontinued. Nodes with new specifications will be added.
- Before scaling out the cluster, check whether its security group configuration is correct. Ensure that an inbound security group rule contains a rule in which **Protocol & Port** is set to **All**, and **Source** is set to a trusted accessible IP address range.

**Step 5** In the **Add Node** dialog box, click **OK**.

**Step 6** A dialog box is displayed in the upper right corner of the page, indicating that the scale-out task is submitted successfully.

The following parameters explain the cluster scale-out process:

- During scale-out: If a cluster is being scaled out, its status is **Scaling out**. The submitted jobs will be executed and you can submit new jobs. You are not allowed to continue to scale out or delete the cluster. You are advised not to restart the cluster or modify the cluster configuration.
- Successful scale-out: The cluster status is **Running**. The resources used in the old nodes and expanded nodes are charged.
- Failed scale-out: The cluster status is **Running**. You can execute jobs and scale out the cluster again.

After the cluster is scaled out, you can view the node information of the cluster on the **Nodes** page.

----End

## Scaling Out a Cluster Billed in Yearly/Monthly Mode

**Step 1** Log in to the MRS console.

**Step 2** Choose **Active Clusters**, select a running cluster, and click its name to switch to the cluster details page.

**Step 3** Click the **Nodes** tab. In the **Operation** column of the node group, click **Scale Out**. The **Scale Out** page is displayed.

The scale-out operation can only be performed on the running clusters.

**Step 4** Set the type of the **System Disk** and **Data Disk**, **Scale-Out Nodes**, **Enable Components** and **Run Bootstrap Action**. The system displays the expiration time of the cluster and the fee required for adding nodes. **Enable Components** and **Run Bootstrap Action** may not be supported by clusters of some versions. The operations are subject to the UI.

### NOTE

- If a bootstrap action is added during cluster creation, the **Run Bootstrap Action** parameter is valid. If this function is enabled, the bootstrap actions added during cluster creation will be run on all the scaled out nodes.
- If the **New Specifications** parameter is available, the specifications that are the same as those of the original nodes have been sold out or discontinued. Nodes with new specifications will be added.
- Before scaling out the cluster, check whether its security group configuration is correct. Ensure that an inbound security group rule contains a rule in which **Protocol & Port** is set to **All**, and **Source** is set to a trusted accessible IP address range.
- Click **Submit Order**.  
On the **Purchase MapReduce Service** page, click **Pay**.
- Click **Confirm order, not pay**.  
On the cluster information page, choose **Fee > My Order** and click **Pay**.

**Step 5** After the payment is successful, return to the MRS management console to view the cluster status.

The following parameters explain the cluster scale-out process:

- During scale-out: If a cluster is being scaled out, its status is **Scaling out**. The submitted jobs will be executed and you can submit new jobs. You are not allowed to continue to scale out or delete the cluster. You are advised not to restart the cluster or modify the cluster configuration.
- Successful scale-out: The cluster status is **Running**. The resources used in the old nodes and expanded nodes are charged.
- Failed scale-out: The cluster status is **Running**. You can execute jobs and scale out the cluster again.

After the cluster is scaled out, you can view the node information of the cluster on the **Nodes** page.

----End

## Adding a Task Node

You can scale out an MRS cluster by manually adding task nodes.

**To add a task node to a custom cluster, perform the following steps:**

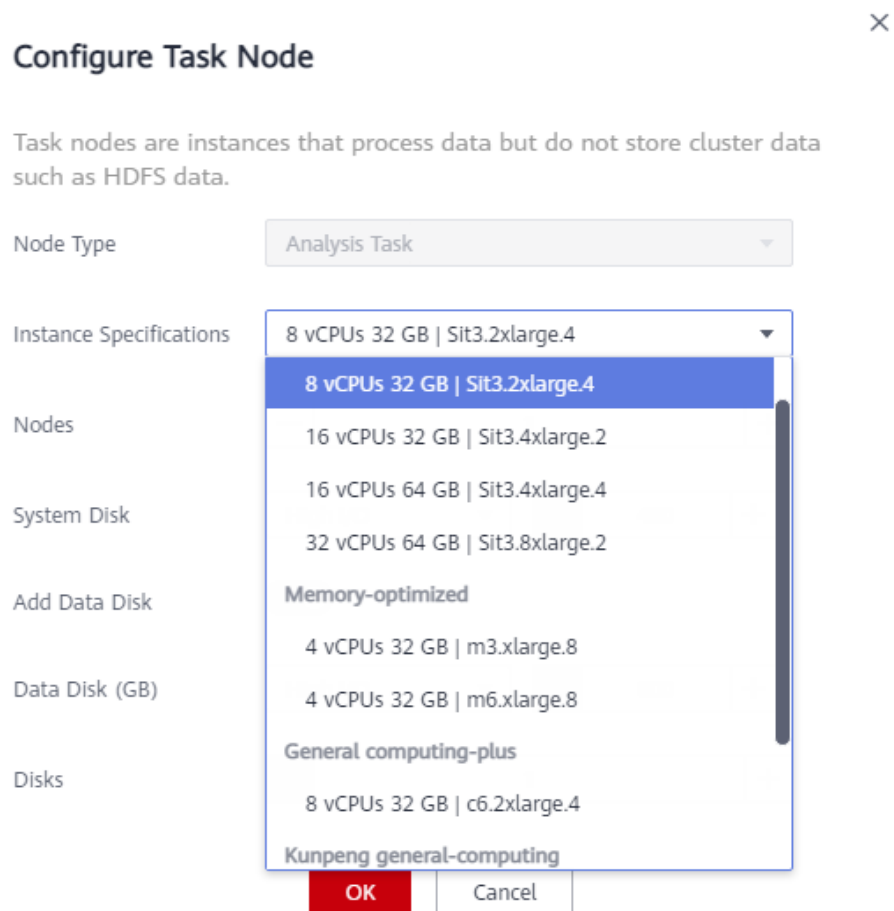
1. On the cluster details page, click the **Nodes** tab and click **Add Node Group**. The **Add Node Group** page is displayed.
2. Select **Task** for **Node Type**. Retain the default value **NM** for **Deploy Roles**. To deploy the NodeManager role, the node type must be **Task**. Set other parameters as required.

**Figure 7-9** Adding a task node group

**To add a task node to a non-custom cluster, perform the following steps:**

1. On the cluster details page, click the **Nodes** tab and click **Configure Task Node**. The **Configure Task Node** page is displayed.

2. On the **Configure Task Node** page, set **Node Type**, **Instance Specifications**, **Nodes**, **System Disk**. In addition, if **Add Data Disk** is enabled, configure the storage type, size, and number of data disks.




3. Click **OK**.

## Adding a Node Group

### NOTE

Used to add node groups and applies to customized clusters of MRS 3.x.

1. On the cluster details page, click the **Nodes** tab and click **Add Node Group**. The **Add Node Group** page is displayed.

**Figure 7-10** Clicking Add Node Group

2. Set the parameters as needed.
3. Click **OK**.

## 7.5.2 Expanding a Data Disk of an MRS Cluster Node

As your business grows, more data disk capacity is needed.

This section describes how to expand a data disk on the MRS console.

If the cluster version is MRS 3.1.0 (patch 3.1.0.0.11 or a later patch must be installed), MRS 3.1.5 (patch 3.1.5.0.3 or a later patch must be installed), or MRS 3.3.0-LTS, expand the data disk capacity by referring to [One-Click Data Disk Expansion](#).

If the cluster version is MRS 3.1.2-LTS.3 or MRS 3.2.0-LTS.1, you can expand the data disk capacity by referring to [Manual Data Disk Expansion](#).

### NOTE

All data disks of selected nodes will be expanded.

## Constraints

- A data disk capacity can only be expanded, and cannot be reduced.
- A data disk can be expanded to a maximum of 32 TB.
- Local disks cannot be expanded.
- System disks cannot be expanded.
- Only the EVS disks, disk partitions, and file systems mounted to the cluster nodes by default during creation can be expanded.
- The cloud server to which an EVS disk is mounted must be in the running state. The EVS disk must be in in-use or available state.

## One-Click Data Disk Expansion

If the cluster version is MRS 3.1.0 (patch 3.1.0.0.11 or a later patch must be installed), MRS 3.1.5 (patch 3.1.5.0.3 or a later patch must be installed), or MRS 3.3.0-LTS, perform the following steps to expand the data disk capacity:

**Step 1** Log in to the MRS console. On the **Active Clusters** page that is displayed, click the name of the desired cluster in the cluster list.

**Step 2** On the MRS details page, click **Nodes**.

**Step 3** Locate the node group where the target disk belongs, click **Expand Data Disk Capacity** in the **Operation** column.

**Step 4** Select the nodes to be expanded, set the target capacity, and click **OK**.

If there are a large number of nodes, you can fuzzy search for nodes by name or IP address, or filter nodes by data disk capacity.

 **NOTE**

- The estimated price is based on the disk capacity displayed and may be inaccurate. You will be billed based on the actual disk capacity of the node. Synchronize the disk information before expanding for more accurate estimated price.
- If the data disk of a node group has been expanded, same expansion will be performed for subsequent scale-out of the node group or cluster cloning.
- When multiple nodes are selected for scale-out, some disks mounted to the nodes may have higher capacity than the target volume. In this case, the disks will not be expanded or billed.
- Only the data disks mounted to a node by default during cluster creation can be expanded. For data disks mounted by yourself, only the EVS disks can be expanded. For details, see [Extending Partitions and File Systems for Data Disks \(Linux\)](#) or contact Huawei Cloud technical support.
- When a disk partition or file system is being expanded, the additional capacity can be expanded only to the tail partition of the disk. A system disk has multiple partitions and cannot be expanded.

**Step 5** Check whether the data disk is expanded.

1. In the disk column, check whether the data disk is expanded to the target volume.
2. Log in to FusionInsight Manager, view the disk information of the expanded node, and check whether the corresponding disk partition is expanded.

----End

## Manual Data Disk Expansion

For clusters of MRS 3.1.2-LTS.3, and MRS 3.2.0-LTS.1 versions, perform the following steps to expand data disks:

1. Contact Huawei Cloud technical support to enable the data disk expansion function.
2. Expand an EVS disk by referring to [One-Click Data Disk Expansion](#).  
Only EVS disks can be expanded to the target volume. Disk partitions and file systems cannot be expanded automatically.
3. Expand a disk partition and file system.  
Download the data disk expansion patches by referring to [Patch Download Addresses](#). Expand a disk partition and file system by referring to the **README.md** file in the patch description.
4. If there is a data disk mounted by a user and the disk partition and file system cannot be expanded, handle the problem by referring to [Extending](#)



[Partitions and File Systems for Data Disks \(Linux\)](#) or contact Huawei Cloud technical support.

## Patch Download Addresses

- **Hong Kong:** [https://mrs-container1-patch-ap-southeast-1.obs.ap-southeast-1.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_Disk\\_Expand\\_Disks\\_Partition\\_Tool\\_Patch.tar.gz](https://mrs-container1-patch-ap-southeast-1.obs.ap-southeast-1.myhuaweicloud.com/MRS_Common_Script/MRS_Disk_Expand_Disks_Partition_Tool_Patch.tar.gz)
- **Singapore:** [https://mrs-container1-patch-ap-southeast-3.obs.ap-southeast-3.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_Disk\\_Expand\\_Disks\\_Partition\\_Tool\\_Patch.tar.gz](https://mrs-container1-patch-ap-southeast-3.obs.ap-southeast-3.myhuaweicloud.com/MRS_Common_Script/MRS_Disk_Expand_Disks_Partition_Tool_Patch.tar.gz)
- **Bangkok:** [https://mrs-container1-patch-ap-southeast-2.obs.ap-southeast-2.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_Disk\\_Expand\\_Disks\\_Partition\\_Tool\\_Patch.tar.gz](https://mrs-container1-patch-ap-southeast-2.obs.ap-southeast-2.myhuaweicloud.com/MRS_Common_Script/MRS_Disk_Expand_Disks_Partition_Tool_Patch.tar.gz)
- **Ulanqab1:** [https://mrs-container1-patch-cn-north-9.obs.cn-north-9.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_Disk\\_Expand\\_Disks\\_Partition\\_Tool\\_Patch.tar.gz](https://mrs-container1-patch-cn-north-9.obs.cn-north-9.myhuaweicloud.com/MRS_Common_Script/MRS_Disk_Expand_Disks_Partition_Tool_Patch.tar.gz)
- **Moscow2:** [https://mrs-container1-patch-ru-northwest-2.obs.ru-northwest-2.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_Disk\\_Expand\\_Disks\\_Partition\\_Tool\\_Patch.tar.gz](https://mrs-container1-patch-ru-northwest-2.obs.ru-northwest-2.myhuaweicloud.com/MRS_Common_Script/MRS_Disk_Expand_Disks_Partition_Tool_Patch.tar.gz)
- **Johannesburg:** [https://mrs-container1-patch-af-south-1.obs.af-south-1.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_Disk\\_Expand\\_Disks\\_Partition\\_Tool\\_Patch.tar.gz](https://mrs-container1-patch-af-south-1.obs.af-south-1.myhuaweicloud.com/MRS_Common_Script/MRS_Disk_Expand_Disks_Partition_Tool_Patch.tar.gz)

## 7.5.3 Scaling In an MRS Cluster

You can reduce the number of core or task nodes to scale in a cluster based on service requirements so that MRS delivers better storage and computing capabilities at lower O&M costs.

The scale-in operation is not allowed for a cluster that is performing active/standby synchronization.

### NOTE

Only pay-per-use clusters can be scaled in. For details about how to scale in a yearly/monthly node, see [Unsubscribing from a Specified Node in a Yearly/Monthly MRS Cluster](#).

## Background

A cluster can have three types of nodes, master, core, and task nodes. Currently, only core and task nodes can be removed. To scale in a cluster, you only need to adjust the number of nodes on the MRS console. MRS then automatically selects the nodes to be removed.

The policies for MRS to automatically select nodes are as follows:

- MRS does not select the nodes with basic components installed, such as ZooKeeper, DBService, KrbServer, and LdapServer, because these basic components are the basis for the cluster to run.
- Core nodes store cluster service data. When scaling in a cluster, ensure that all data on the core nodes to be removed has been migrated to other nodes. You can perform follow-up scale-in operations only after all component services

are decommissioned, for example, removing nodes from Manager and deleting ECSs. When selecting core nodes, MRS preferentially selects the nodes with a small amount of data and healthy instances to be decommissioned to prevent decommissioning failures. For example, if DataNodes are installed on core nodes in an analysis cluster, MRS preferentially selects the nodes with small data volume and good health status during scale-in.

When core nodes are removed, their data is migrated to other nodes. If the user business has cached the data storage path, the client will automatically update the path, which may increase the service processing latency temporarily. Cluster scale-in may slow the response of the first access to some HBase on HDFS data. You can restart HBase or disable or enable related tables to resolve this issue.

- Task nodes are computing nodes and do not store cluster data. Data migration is not involved in removing task nodes. Therefore, when selecting task nodes, MRS preferentially selects nodes whose health status is faulty, unknown, or subhealthy. On the **Components** tab of the MRS console, click a service and then the **Instances** tab to view the health status of the node instances.

## Scale-In Verification Policy

To prevent component decommissioning failures, components provide different decommissioning constraints. Scale-in is allowed only when the constraints of all installed components are met. [Table 7-8](#) describes the scale-in verification policies.

**Table 7-8** Decommissioning constraints

| Component          | Constraint                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HDFS/DataNode      | <p>The number of available nodes after the scale-in is greater than or equal to the number of HDFS copies and the total HDFS data volume does not exceed 80% of the total HDFS cluster capacity.</p> <p>This ensures that the remaining space is sufficient for storing existing data after the scale-in and reserves some space for future use.</p> <p><b>NOTE</b></p> <p>To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total.</p> |
| HBase/RegionServer | <p>The total available memory of RegionServers on all nodes except the nodes to be removed is greater than 1.2 times of the memory which is currently used by RegionServers on these nodes.</p> <p>This ensures that the node to which the region on a decommissioned node is migrated has sufficient memory to bear the region of the decommissioned node.</p>                                                                                                                                                           |

| Component                       | Constraint                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storm/<br>Supervisor            | After the scale-in, ensure that the number of slots in the cluster is sufficient for running the submitted tasks.<br>This prevents no sufficient resources being available for running the stream processing tasks after the scale-in.                                                                                                                                                                                                                        |
| Flume/<br>FlumeServer           | If FlumeServer is installed on a node and Flume tasks have been configured for the node, the node cannot be deleted.<br>This prevents the deployed service program from being deleted by mistake.                                                                                                                                                                                                                                                             |
| ClickHouse/<br>ClickHouseServer | For details, see <a href="#">Constraints on ClickHouseServer Scale-in</a> .<br>This ensures that data on the decommissioned nodes is migrated to in-use nodes.                                                                                                                                                                                                                                                                                                |
| Kudu/<br>KuduTserver            | Rule: When you decommission a Kudu Tserver, all other Kudu instance nodes in the cluster must be normal for the process to succeed.<br>Cause: During the decommissioning of a KuduTserver, a rebalance command is executed to migrate the tablets from the decommissioned instance to other KuduTserver nodes in the cluster. If any other KuduTserver nodes are in an abnormal state, the rebalance command will fail, resulting in decommissioning failure. |

## Scaling In a Cluster by Specifying the Node Quantity

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** Click the **Nodes** tab. In the **Operation** column of the node group, click **Scale In** to go to the **Scale In** page.

This operation can be performed only when the cluster and all nodes in it are running.

**Step 4** Set **Scale-In Type** to **Node quantity**.

**Step 5** Set **Scale-In Nodes** and click **OK**.

## Scale In

To improve scale-in reliability, MRS features standard scale-in rules for big data service components. If you perform the scale-in, the server and data disks will be deleted and cannot be recovered. [Learn more](#)

|                                                                                              |                                                                                                  |
|----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Scale-In Type                                                                                | <input checked="" type="radio"/> Node quantity <input type="radio"/> Specific node               |
| Node Type                                                                                    | Analysis Core                                                                                    |
| Current Nodes                                                                                | 4                                                                                                |
| * Scale-In Nodes                                                                             | <input type="button" value="-"/> <input type="text" value="1"/> <input type="button" value="+"/> |
| <input type="checkbox"/> I understand the consequences of performing the scale-in operation. |                                                                                                  |
| <input type="button" value="OK"/> <input type="button" value="Cancel"/>                      |                                                                                                  |

### NOTE

- Before scaling in the cluster, check whether its security group configuration is correct. Ensure that an inbound security group rule contains a rule in which **Protocol & Port** is set to **All**, and **Source** is set to a trusted accessible IP address range.
- If damaged data blocks exist in HDFS, the cluster may fail to be scaled in. Contact Huawei Cloud technical support.

**Step 6** A dialog box displayed in the upper right corner of the page indicates that the task of removing the node is submitted successfully.

The cluster scale-in process is explained as follows:

- During scale-in: The cluster status is **Scaling In**. The submitted jobs will be executed, and you can submit new jobs. You are not allowed to continue to scale in or delete the cluster. You are advised not to restart the cluster or modify the cluster configuration.
- Successful scale-in: The cluster status is **Running**. The resources used after the cluster scale-in are billed.
- Failed scale-in: The cluster status is **Running**. You can execute jobs or scale-in the cluster again.

After the cluster is scaled in, you can view the node information of the cluster on the **Nodes** page.

----End

## Scaling In a Cluster by Removing Nodes that Are No Longer Needed

Delete a node when it is no longer needed. Before deleting a node, decommission the role instance of the component and back up the node's data. For details about

how to remove ClickHouseServer nodes, see [Scaling In ClickHouseServer Nodes](#). Only pay-per-use nodes can be scaled in.

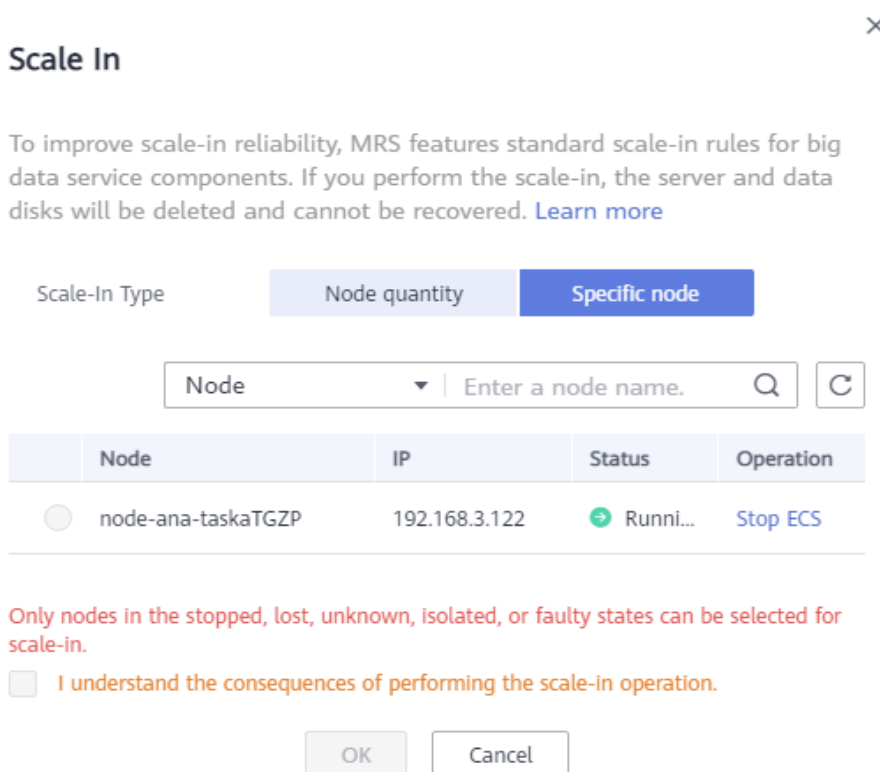
- Step 1** Log in to the MRS management console.
- Step 2** Click the name of the cluster to go to its details page.
- Step 3** Click the **Nodes** tab.
- Step 4** Locate the row that contains the target node group and click **Scale In** in the **Operation** column to go to the **Scale In** page.
- Step 5** Set **Scale-In Type** to **Specific node** and select the node to be removed.

Nodes in the **Stopped**, **Lost**, **Unknown**, **Isolated**, or **Faulty** status can be specified for scale-in. If the node cannot be selected, click **Stop ECS** to go to the ECS console to stop the node. On the cluster details page of the MRS console, click the **Alarms** tab and check whether any service fault alarms are generated after the node is stopped. If no such an alarm is generated, go back to the **Scale In** page and select the corresponding node for scale-in. If such an alarm is generated, clear the alarm before the scale-in.

**NOTE**

This operation may cause data loss. Decommission nodes before scale-in. For details, see [Decommissioning and Recommissioning an MRS Role Instance](#).

**Figure 7-11** Removing a specific node



- Step 6** Select **I understand the consequences of performing the scale-in operation**, Click **OK**.

**Step 7** Click the **Components** tab and check whether each component is normal. If any component is abnormal, wait for 5 to 10 minutes and check the component status again. If the fault persists, contact Huawei Cloud technical support.

**Step 8** Click the **Alarms** tab and check whether there are exception alarms. If there are exception alarms, clear them before performing other operations.

----End

## 7.5.4 Scaling In ClickHouseServer Nodes

If ClickHouse is deployed in the MRS cluster, check the data to prevent data loss during node deletion before scaling in ClickHouseServer nodes.

### Constraints on ClickHouseServer Scale-in

**Table 7-9** Scale-in constraints

| Dimension | Constraints                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scale     | <ul style="list-style-type: none"><li>If a cluster has only one shard, the instance nodes cannot be removed from the cluster.</li><li>Multiple instance nodes in the same shard <b>must be decommissioned or recommissioned at the same time</b>. To query cluster shard information, perform the following steps:<ol style="list-style-type: none"><li>Log in to the node where the HDFS clients are installed as the client installation user and run the following commands:<pre>cd <i>Client installation directory</i> source bigdata_env Security mode kinit <i>ClickHouse service user</i> clickhouse client --host <i>IP address of the ClickHouse instance</i> --port 9440 --secure</pre>In normal mode, run the following command:<pre>clickhouse client --host <i>IP address of the ClickHouse instance</i> --user <i>Username</i> --password --port 9000</pre><i>Enter the user password.</i></li><li>Run the following command to query the cluster shard information:<pre>select cluster,shard_num,replica_num,host_name from system.clusters;</pre></li></ol></li></ul> |

| Dimension       | Constraints                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Storage | <p>Ensure that the disk space of nodes that will not be decommissioned is sufficient for storing data of all decommissioned nodes. There must be approximately 10% redundant storage space after decommissioning to ensure that the remaining instances can run properly. The procedure is as follows:</p> <ol style="list-style-type: none"><li>1. Run the following command to check the disk usage on each node:<br/><b>select * from system.disks;</b><br/><b>free_space</b> indicates the free disk space, and <b>total_space</b> indicates the total disk space. The used space is calculated by subtracting the value of <b>free_space</b> from that of <b>total_space</b>, and its unit is byte.</li><li>2. Run the preceding command on a node you want to decommission and calculate the data volume on the node using the preceding formula.</li><li>3. Run the preceding command on a node that will not be decommissioned, and then use the following formula: (Value of <b>free_space</b> - Data volume of the node to be decommissioned)/Value of <b>total_space</b>. If the result is greater than 10%, the node can be decommissioned.</li></ol> |
| Cluster Status  | <p>If there is any faulty ClickHouseServer instance node in the cluster, all instance nodes in the cluster cannot be decommissioned.</p> <p>Log in to Manager, choose <b>Cluster &gt; Services &gt; ClickHouse</b>, click <b>Instance</b>, and view the running status of each node in the cluster.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Database        | <p>If a database is deployed only on an instance node you want to decommission, the instance node cannot be decommissioned. To remove the instance node, you need to create the database on all ClickHouseServer instance nodes in the cluster. The procedure is as follows:</p> <ol style="list-style-type: none"><li>1. Run the <b>select * from system.databases;</b> command to collect the database list of each node.<br/><b>name</b> indicates the database name. <b>engine</b> indicates the database engine, and the default value is <b>Atomic</b>. If the default engine is used, you do not need to specify the engine when creating a table.</li><li>2. For the database deployed only on the instance node to be decommissioned, run the following command to create the database:<br/><b>create database xxx engine=xxx on cluster xxx;</b></li></ol>                                                                                                                                                                                                                                                                                              |

| Dimension                  | Constraints                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Non-replicated Table | <p>If a local non-replicated table is deployed only on an instance node you want to decommission, the instance node cannot be decommissioned. To decommission the node, create a local non-replicated table with the same name on any node that will not be decommissioned.</p> <p>For example, the current cluster has two shards, shard 1 has two nodes A and B, and shard 2 has two nodes C and D. The non-replicated table <b>test</b> was created without the <b>ON CLUSTER</b> keyword, so the table is created only on node A.</p> <p>In this case, to decommission nodes A and B in shard 1, you need to create the table <b>test</b> on node C or D in shard 2.</p> <p>Run the following command to list the data tables of each node:</p> <pre><b>select database,name,engine,create_table_query from system.tables where database != 'system';</b></pre> <p>Perform the following operations according to the result:</p> <ul style="list-style-type: none"> <li>• Check the <b>engine</b> column. The table that does not contain the <b>Replicated</b> field is a local non-replicated table.</li> <li>• If there are no replicated tables on any nodes that will not be decommissioned, create one based the table created by <b>create_table_query</b>. The following creation statement is an example:<br/><b>CREATE TABLE {database}.{table} ('column name' type...) ENGINE = MergeTree;</b></li> </ul> |
| Replicated Table           | <p>If a replicated table exists only on some nodes in the cluster, the nodes where the replicated table is deployed cannot be decommissioned. You need to manually create the replicated table on all instance nodes where no replicated table is deployed in the cluster before decommissioning.</p> <p>For example, the current cluster has two shards, shard 1 has two nodes A and B, and shard 2 has two nodes C and D. The replicated table <b>test</b> was created without the <b>ON CLUSTER</b> keyword, so the table is created only on nodes A and B.</p> <p>To decommission nodes A and B in shard 1, you need to create the table <b>test</b> on nodes C and D in shard 2.</p> <p>Run the following command to list the data tables of each node:</p> <pre><b>select database,name,engine,create_table_query from system.tables where database != 'system';</b></pre> <p>Perform the following operations according to the result:</p> <ul style="list-style-type: none"> <li>• Check the <b>engine</b> column. The table that contains the <b>Replicated</b> field is a replicated table.</li> <li>• If there are no replicated tables on any nodes that will not be decommissioned, create one based the table created by <b>create_table_query</b>.</li> </ul>                                                                                                                                             |



| Dimension         | Constraints                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Distributed Table | <p>Distributed tables will not be migrated automatically for decommissioning. Create distributed tables on the nodes that will not be decommissioned.</p> <p>Run the following command to list data tables of each node and check the <b>engine</b> column. These tables are distributed tables if this column contains field <b>Distributed</b>.</p> <p><b>select database,name,engine from system.tables where database != 'system';</b></p> <p><b>NOTE</b><br/>Creating distributed tables on these nodes will not affect the decommissioning, but may affect subsequent service operations.</p> |
| View              | <p>Views will not be automatically migrated for decommissioning, and views do not store data. Run the following command to list data tables of each node and check the <b>engine</b> column. These tables are views if this column contains field <b>View</b>.</p> <p><b>select database,name,engine from system.tables where database != 'system';</b></p> <p>Run the following command to delete the views one by one:</p> <p><b>drop view {database_name}.{table_name};</b></p>                                                                                                                  |

| Dimension          | Constraints                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Materialized Views | <p>Materialized views will not be automatically migrated for Decommissioning. Create materialized views on the nodes that will not be decommissioned. If the materialized view of a node to be decommissioned does not display the specified aggregation table but uses an embedded table, the node cannot be decommissioned.</p> <p>Run the following command to list data tables of each node and check the <b>engine</b> column. These tables are materialized views if this column contains field <b>MaterializedView</b>.</p> <pre><b>select database,name,engine, create_table_query from system.tables where database != 'system';</b></pre> <p>Embedded tables are initialized with a <b>POPULATE</b> field in their <b>create_table_query</b> column. Views are created at initialization, and new data is ignored. Aggregation tables without the <b>POPULATE</b> field insert new data directly into views and support tables, requiring manual data loading. The table creation operations of the aggregation table and embedded table are different.</p> <p>Perform the following operations to process the materialized views of the node to be decommissioned:</p> <ol style="list-style-type: none"> <li>1. Record the materialized views and delete them.<br/><b>drop view {database_name}.{table_name};</b></li> <li>2. After the node decommissioning is complete, delete and recreate the corresponding materialized views on in-use nodes to update the materialized views.</li> <li>3. To create an aggregation table, specify <b>WHERE</b> to search for historical data and manually import the historical data to the materialized views. Otherwise, historical data cannot be imported to the materialized views based on unified conditions. As a result, data is imported repeatedly. For example, an update point can be specified to ensure that data before the update point is manually loaded in <b>INSERT</b> mode. <ul style="list-style-type: none"> <li>• Add <b>WHERE { Time field (for example, date)}&gt;= toDate ({ Current time (for example, '2022-12-01 00:00:00')})</b> to the table creation statement.</li> <li>• <b>insert into {table} select {Table field} from {Source table} where {Time field}&lt; toDate ({Current time})</b> is used to load original data.</li> </ul> </li> <li>4. Embedded tables will lose data generated during table creation. You can specify <b>WHERE</b> to filter out all historical data. In this case, an empty table is created, and you only need to manually insert all data in the historical data source table.</li> </ol> |

| Dimension                     | Constraints                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tables of Third-Party Engines | <p>Currently, tables of third-party engines cannot be automatically migrated for decommissioning.</p> <p>Run the following command to list data tables of each node and check the <b>engine</b> column. These tables are tables of third-party engines if this column does not contain any of the following fields: <b>MergeTree</b>, <b>View</b>, <b>MaterializedView</b>, <b>Distributed</b>, and <b>Log</b>. (The <b>engine</b> column of a third-party engine table may contain field <b>Memory</b>, <b>HDFS</b>, or <b>MySQL</b>.)</p> <pre><b>select database,name,engine from system.tables where database != 'system';</b></pre> <p>Create third-party engine tables on the nodes that will not be decommissioned and delete those from the nodes that will be decommissioned.</p>                                                                                                                                                                                                                                                                                                                            |
| Detached Data                 | <p>If the table on a node to be decommissioned has been detached and data still exists in the <b>detached</b> directory, the node cannot be decommissioned. You need to attach the data in the <b>detached</b> directory to other directories before decommissioning.</p> <ol style="list-style-type: none"> <li>1. Run the following command to view the <b>system.detached_parts</b> system catalog of the node to be decommissioned:<br/><pre><b>select * from system.detached_parts;</b></pre></li> <li>2. If <b>detached part</b> data exists and these partitions are no longer used, run the following command to delete the <b>detached part</b> data:<br/><pre><b>ALTER TABLE {table_name} DROP DETACHED PARTITION {partition_expr} SETTINGS allow_drop_detached = 1;</b></pre></li> <li>3. Run the following command to check whether there is any <b>detached part</b> data in the <b>system.detached_parts</b> system catalog:<br/><pre><b>select * from system.detached_parts;</b></pre></li> </ol> <p>If the command output is empty, there is no <b>detached part</b> data in this system catalog.</p> |

## Scaling In ClickHouseServer Nodes

Before removing ClickHouseServer instance nodes, you need to decommission them. Multiple node replicas of the same shard **must be decommissioned at the same time**. If there is a faulty ClickHouseServer instance node in the cluster, all instance nodes of the cluster cannot be decommissioned. For more constraints, see [Constraints on ClickHouseServer Scale-in](#).

 NOTE

- Perform the decommissioning in idle hours because the operation will occupy certain bandwidth resources.
- The decommissioning operation can be performed only to ClickHouseServer. ClickHouseBalancer cannot be decommissioned.
- **This operation is only supported for MRS 3.1.2 and later.**

**Step 1** Use PuTTY to log in to the node where ClickHouseServer is installed as user **root** and run the following command:

```
echo 'select * from system.clusters' | curl -k 'https://IP address of the node
where the ClickHouseServer instance is located:Port number/' -u ck_user.Password
--data-binary @-
```

Record the nodes of the same shard. In the following command output, the nodes with the same number in bold belong to the same shard.

```
[root@kwephispra44948 ~]# echo 'select * from system.clusters' | curl -k 'https://10.112.17.189:21422/' -u
ck_user:Bigdata_2013 --data-binary @-
default_cluster 1 1 1 kwephispra44947 10.112.17.150 21427 0 0 0
default_cluster 1 1 2 kwephispra44948 10.112.17.189 21427 0 0 0
```

 NOTE

- To view the port number of ClickHouseServer instance nodes, log in to FusionInsight Manager, choose **Cluster > Services > ClickHouse**, click **Configuration > All Configurations**, and choose **ClickHouseServer (Role)** on the left.  
In security mode (Kerberos authentication enabled), check the value of **https\_port**, which is the port of a ClickHouseServer instance node.  
In common mode (Kerberos authentication disabled), check the value of **http\_port**, which is the port of a ClickHouseServer instance node.
- **ck\_user** indicates the created ClickHouse user, which must be bound to a role with the ClickHouse administrator permission. For details about how to create a user and a role, see [Creating an MRS Cluster User](#) and [Managing MRS Cluster Roles](#), respectively.

**Step 2** Log in to the MRS console and click the cluster name to go to the cluster details page.

**Step 3** Click the **Components** tab and click **ClickHouse**. Then switch to **Instances**, select the **ClickHouseServer** instances to be removed, click **More**, and select **Decommission**.

**Step 4** Click the **Components** tab and click **ClickHouse**. Then click **More**, and select **Synchronize Configuration**.

**Step 5** Click the **Nodes** tab and click the ClickHouseServer instance node that has been decommissioned.

**Step 6** On the ECS page, click **Stop**. In the displayed dialog box, select **Forcibly stop the preceding ECSs** and click **Yes**.

**Step 7** Go back to the MRS console, click the **Nodes** tab, locate the row that contains the target node group, and click **Scale In** in the **Operation** column to go to the **Scale In** page.

**Step 8** Set **Scale-In Type** to **Specific node** and select the node to be removed.

**Step 9** Select **I understand the consequences of performing the scale-in operation**. Click **OK**.

- Step 10** Click the **Components** tab and check whether each component is normal. If any component is abnormal, wait for 5 to 10 minutes and check the component status again. If the fault persists, contact Huawei Cloud technical support.
- Step 11** Click the **Alarms** tab and check whether there are exception alarms. If there are exception alarms, clear them before performing other operations.

----End

## 7.5.5 Unsubscribing from a Specified Node in a Yearly/ Monthly MRS Cluster

You can reduce the number of specific nodes to scale in a cluster so that MRS delivers better storage and computing capabilities at lower O&M costs based on service requirements.

Currently, you can unsubscribe from a maximum of 20 Core nodes at a time, but there must be at least 2 Core nodes available after unsubscription.

### NOTE

You can unsubscribe from a node only after the node is successfully isolated or decommissioned. Otherwise, data loss may occur.

## Usage Restrictions

- If the number of Core nodes in the cluster is less than or equal to the number of HDFS copies, MRS does not support node unsubscription to ensure data reliability. The number of HDFS copies can be queried using the **dfs.replication** parameter in the HDFS parameter configuration.
- MRS does not support unsubscription from nodes where ZooKeeper, Kudu, Kafka, or ClickHouse is deployed.

## How to Unsubscribe from a Specified Node in a Yearly/Monthly Cluster

- Step 1** Disable the auto-renewal function of the cluster where the node to be unsubscribed from is located. For details, see [Disabling Auto-Renewal](#).
- Step 2** Log in to the MRS console.
- Step 3** On the **Active Clusters** page, and click the name of the target cluster to be operated. The cluster details page is displayed.
- Step 4** On the **Dashboard** page of the cluster, click **Synchronize** on the right of **IAM User Sync**.
- Step 5** Unsubscribe from or isolate a node.

 NOTE

Currently, only clusters of the following versions support unsubscription from specified nodes in yearly/monthly clusters. For clusters of other versions, contact technical support.

- MRS 2.1.0 (patch 2.1.0.5 or later)
- MRS 3.1.0 (patch 3.1.0.0.2 or later)
- MRS 3.1.5
- MRS 3.2.0-LTS.1 (patch 3.2.0-LTS.1.3 or later)
- If the cluster is earlier than MRS 2.x:
  - a. Click **Isolate Node** in the **Operation** column of the node group to be unsubscribed from.
  - b. Select the node to be unsubscribed from and click **OK**.


The time required for isolating a node depends on the data volume on the node. A larger data volume indicates a longer time.

After the node is isolated, the node status changes to **Isolated**. The **Unsubscribe from Node** button is displayed on the **Nodes** tab page.
- If the cluster is MRS 3.1.0, 3.1.5, or 3.2.0-LTS.1:
  - a. Click **Decommission Node** in the **Operation** column of the node group to be unsubscribed from.
  - b. Select the node to be decommissioned and click **OK**.

The time required for decommissioning a node depends on the data volume on the node. A larger data volume indicates a longer time.

After the node is decommissioned, the node status changes to **Decommissioned**.

 NOTE


- For clusters of MRS 3.x (except MRS 3.1.0, 3.1.5, and 3.2.0-LTS.1), perform the following operations:
  1. To unsubscribe a node that has HDFS DataNode, YARN NodeManager, or HBase RegionServer instances, decommission these instances on the Manager first. For details, see [Decommissioning and Recommissioning an MRS Role Instance](#).
  2. Log in to the MRS management console, on the **Nodes** tab, select the nodes to be unsubscribed from, and choose **Node Operation** > **Isolate Host**. After the host is isolated, contact the technical engineer to unsubscribe from the node.
- Only one node can be isolated or decommissioned at a time. You can unsubscribe from a node only after the node is successfully isolated or decommissioned.
- If the node fails to isolate or decommission, log in to Manager. Click , search for the name of the task that fails to isolate or decommission the host in the task list, click the name, and rectify the fault as prompted.

**Step 6** On the cluster details page, choose **Nodes** > **Unsubscribe from Node**.

**Step 7** Select the node to be unsubscribed from and click **OK**.

Currently, you can unsubscribe from a maximum of 20 Core nodes at a time, but there must be at least 2 Core nodes available after unsubscription.

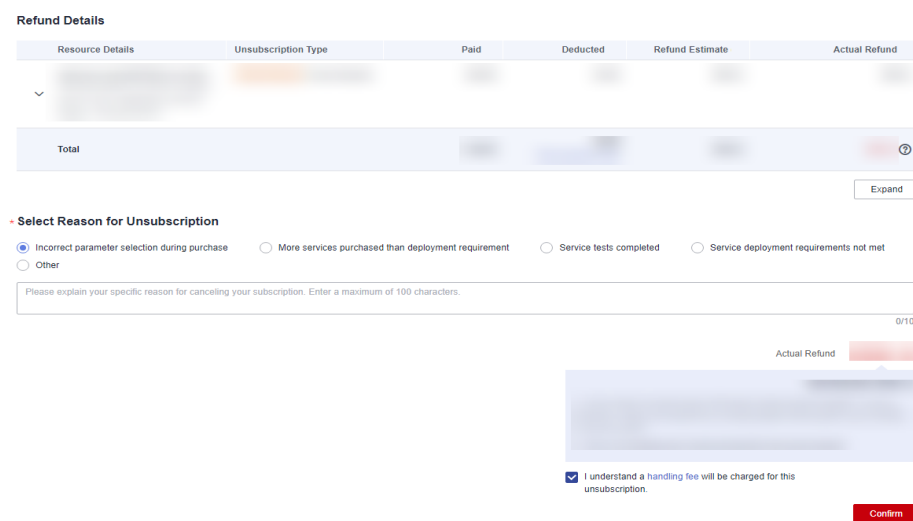
**Figure 7-12** Selecting a node to be unsubscribed from



**Step 8** On the **Refund Details** page, select "I understand a handling fee will be charged for this unsubscription" as prompted if needed, and click **Confirm**. If this check box does not appear, skip this step.

The cluster status changes to **Scaling in**. After the scale-in is complete, the cluster status changes to **Running**, and the specified node is deleted.

**Figure 7-13** Refund details



**Step 9** (Optional) To enable auto-renewal for a cluster, see [Enabling Auto-Renewal](#).

----End

## 7.5.6 MRS Task Node Auto Scaling

### 7.5.6.1 Automatic Scaling of Task Nodes in an MRS Cluster

In big data application scenarios, especially real-time data analysis and processing, the number of cluster nodes needs to be dynamically adjusted according to data volume changes to provide the required number of resources. The auto scaling function of MRS enables the task nodes of a cluster to be automatically scaled to match cluster loads. If the data volume changes periodically, you can configure an

auto scaling rule so that the number of task nodes can be automatically adjusted in a fixed period of time before the data volume changes.

- Auto scaling rules: You can increase or decrease task nodes based on real-time cluster loads. Auto scaling will be triggered with a certain delay when the data volume changes.
- Resource plans: Set the task node quantity based on the time range. If the data volume changes periodically, you can create resource plans to resize the cluster before the data volume changes, thereby avoiding delays in increasing or decreasing resources.

You can configure either auto scaling rules or resource plans or both to trigger auto scaling. Configuring both resource plans and auto scaling rules improves the cluster node scalability to cope with occasionally unexpected data volume peaks.

In some service scenarios, resources need to be reallocated or service logic needs to be modified after cluster scale-out or scale-in. If you manually scale out or scale in a cluster, you can log in to cluster nodes to reallocate resources or modify service logic. If you use auto scaling, MRS enables you to customize automation scripts for resource reallocation and service logic modification. Automation scripts can be executed before and after auto scaling and automatically adapt to service load changes, all of which eliminates manual operations. In addition, automation scripts can be fully customized and executed at various moments, meeting your personalized requirements and improving auto scaling flexibility.

- Auto scaling rules:
  - You can set a maximum of five rules for scaling out or in a cluster, respectively.
  - The system determines the scale-out and then scale-in based on your configuration sequence. Important policies take precedence over other policies to prevent repeated triggering when the expected effect cannot be achieved after a scale-out or scale-in.
  - Comparison factors include greater than, greater than or equal to, less than, and less than or equal to.
  - Cluster scale-out or scale-in can be triggered only after the configured metric threshold is reached for consecutive  $5n$  (the default value of  $n$  is 1) minutes.
  - After each scale-out or scale-in, there is a cooling duration that is greater than 0 and lasts 10 minutes by defaults.
  - In each cluster scale-out or scale-in, at least one node and at most 100 nodes can be added or reduced.
  - The number of task nodes in a cluster is limited to the default number of nodes configured by users or the node quantity range in the resource plan that takes effect in the current time period. The node quantity range in the resource plan that takes effect in the current time period has a higher priority.
- Resource plans (setting the number of Task nodes by time range):
  - You can specify a Task node range (minimum number to maximum number) in a time range. If the number of Task nodes is beyond the Task node range in a resource plan, the system triggers cluster scale-out or scale-in.



- You can set a maximum of five resource plans for a cluster.
- A resource plan cycle is by day. The start time and end time can be set to any time point between 00:00 and 23:59. The start time must be at least 30 minutes earlier than the end time. Time ranges configured for different resource plans cannot overlap.
- After a resource plan triggers cluster scale-out or scale-in, there is 10-minute cooling duration. Auto scaling will not be triggered again within the cooling time.
- When a resource plan is enabled, the number of Task nodes in the cluster is limited to the default node range configured by you in other time periods except the time period configured in the resource plan.
- Automation scripts:
  - You can set an automation script so that it can automatically run on cluster nodes when auto scaling is triggered.
  - You can set a maximum number of 10 automation scripts for a cluster.
  - You can specify an automation script to be executed on one or more types of nodes.
  - Automation scripts can be executed before or after scale-out or scale-in.
  - Before using automation scripts, upload them to a cluster VM or OBS file system in the same region as the cluster. The automation scripts uploaded to the cluster VM can be executed only on the existing nodes. If you want to make the automation scripts run on the new nodes, upload them to the OBS file system.

## Node Auto Scaling Metrics

- **Node group dimension policy**

When adding a rule, you can refer to [Table 7-10](#) to configure the corresponding metrics.

**Table 7-10** Auto scaling metrics

| Cluster Type      | Metric                       | Value Type | Description                                                                                                  |
|-------------------|------------------------------|------------|--------------------------------------------------------------------------------------------------------------|
| Streaming cluster | StormSlotAvailable           | Integer    | Number of available Storm slots<br>Value range: 0 to 2147483646                                              |
|                   | StormSlotAvailablePercentage | Percentage | Percentage of available Storm slots, that is, the proportion of the available slots<br>Value range: 0 to 100 |
|                   | StormSlotUsed                | Integer    | Number of the used Storm slots<br>Value range: 0 to 2147483646                                               |
|                   | StormSlotUsedPercentage      | Percentage | Percentage of the used Storm slots, that is, the proportion of the used slots<br>Value range: 0 to 100       |

| Cluster Type     | Metric                                   | Value Type | Description                                                                                                                                  |
|------------------|------------------------------------------|------------|----------------------------------------------------------------------------------------------------------------------------------------------|
|                  | StormSupervisorMemAverageUsage           | Integer    | Average memory usage of the Supervisor process of Storm<br>Value range: 0 to 2147483646                                                      |
|                  | StormSupervisorMemAverageUsagePercentage | Percentage | Average percentage of the memory used by the Supervisor process of Storm to the total memory of the system<br>Value range: 0 to 100          |
|                  | StormSupervisorCPUAverageUsagePercentage | Percentage | Average percentage of the CPUs used by the Supervisor process of Storm to the total CPUs<br>Value range: 0 to 6000                           |
| Analysis cluster | YARNAppPending                           | Integer    | Number of pending tasks on YARN<br>Value range: 0 to 2147483646                                                                              |
|                  | YARNAppPendingRatio                      | Ratio      | Ratio of pending tasks on YARN, that is, the ratio of pending tasks to running tasks on YARN<br>Value range: 0 to 2147483646                 |
|                  | YARNAppRunning                           | Integer    | Number of running tasks on YARN<br>Value range: 0 to 2147483646                                                                              |
|                  | YARNContainerAllocated                   | Integer    | Number of containers allocated to YARN<br>Value range: 0 to 2147483646                                                                       |
|                  | YARNContainerPending                     | Integer    | Number of pending containers on YARN<br>Value range: 0 to 2147483646                                                                         |
|                  | YARNContainerPendingRatio                | Ratio      | Ratio of pending containers on YARN, that is, the ratio of pending containers to running containers on YARN.<br>Value range: 0 to 2147483646 |
|                  | YARNCPUAllocated                         | Integer    | Number of virtual CPUs (vCPUs) allocated to YARN<br>Value range: 0 to 2147483646                                                             |
|                  | YARNCPUAvailable                         | Integer    | Number of available vCPUs on YARN<br>Value range: 0 to 2147483646                                                                            |

| Cluster Type | Metric                        | Value Type | Description                                                                                                                          |
|--------------|-------------------------------|------------|--------------------------------------------------------------------------------------------------------------------------------------|
|              | YARNCPUAvailablePercentage    | Percentage | Percentage of available vCPUs on YARN, that is, the proportion of available vCPUs to total vCPUs<br>Value range: 0 to 100            |
|              | YARNCPUPending                | Integer    | Number of pending vCPUs on YARN<br>Value range: 0 to 2147483646                                                                      |
|              | YARNMemoryAllocated           | Integer    | Memory allocated to YARN. The unit is MB.<br>Value range: 0 to 2147483646                                                            |
|              | YARNMemoryAvailable           | Integer    | Available memory on YARN. The unit is MB.<br>Value range: 0 to 2147483646                                                            |
|              | YARNMemoryAvailablePercentage | Percentage | Percentage of available memory on YARN, that is, the proportion of available memory to total memory on YARN<br>Value range: 0 to 100 |
|              | YARNMemoryPending             | Integer    | Pending memory on YARN<br>Value range: 0 to 2147483646                                                                               |

 **NOTE**

- When the value type is percentage or ratio in [Table 7-10](#), the valid value can be accurate to percentile. The percentage metric value is a decimal value with a percent sign (%) removed. For example, 16.80 represents 16.80%.
- Hybrid clusters support all metrics of analysis and streaming clusters.

- **Resource pool policy**

When adding a rule, you can refer to [Table 7-11](#) to configure the corresponding metrics.

 **NOTE**

Auto scaling policies can be configured for a cluster by resource pool in MRS 3.1.5 or later.

**Table 7-11** Rule configuration description

| Cluster Type            | Metric                                | Value Type | Description                                                                                                                                               |
|-------------------------|---------------------------------------|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Analysis/Custom cluster | ResourcePoolMemoryAvailable           | Integer    | Available memory on YARN in the resource pool. The unit is MB.<br>Value range: 0 to 2147483646                                                            |
|                         | ResourcePoolMemoryAvailablePercentage | Percentage | Percentage of available memory on YARN in the resource pool, that is, the proportion of available memory to total memory on YARN<br>Value range: 0 to 100 |
|                         | ResourcePoolCPUAvailable              | Integer    | Number of available vCPUs on YARN in the resource pool<br>Value range: 0 to 2147483646                                                                    |
|                         | ResourcePoolCPUAvailablePercentage    | Percentage | Percentage of available vCPUs on YARN in the resource pool. that is, the proportion of available vCPUs to total vCPUs<br>Value range: 0 to 100            |

When adding a resource plan, you can set parameters by referring to [Table 7-12](#).

**Table 7-12** Configuration items of a resource plan

| Configuration Item | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Effective On       | The effective date of a resource plan. <b>Daily</b> is selected by default. You can also select one or multiple days from Monday to Sunday.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Time Range         | Start time and End time of a resource plan are accurate to minutes, with the value ranging from <b>00:00</b> to <b>23:59</b> . For example, if a resource plan starts at 8:00 and ends at 10:00, set this parameter to 8:00-10:00. The end time must be at least 30 minutes later than the start time.                                                                                                                                                                                                                                                                                                     |
| Node Range         | The number of nodes in a resource plan ranges from <b>0</b> to <b>500</b> . In the time range specified in the resource plan, if the number of Task nodes is less than the specified minimum number of nodes, it will be increased to the specified minimum value of the node range at a time. If the number of Task nodes is greater than the maximum number of nodes specified in the resource plan, the auto scaling function reduces the number of Task nodes to the maximum value of the node range at a time. The minimum number of nodes must be less than or equal to the maximum number of nodes. |

 **NOTE**

- When a resource plan is enabled, the **Default Range** value on the auto scaling page forcibly takes effect beyond the time range specified in the resource plan. For example, if **Default Range** is set to **1-2**, **Time Range** is between **08:00-10:00**, and **Node Range** is **4-5** in a resource plan, the number of Task nodes in other periods (0:00-8:00 and 10:00-23:59) of a day is forcibly limited to the default node range (1 to 2). If the number of nodes is greater than 2, auto scale-in is triggered; if the number of nodes is less than 1, auto scale-out is triggered.
  - When a resource plan is not enabled, the **Default Range** takes effect in all time ranges. If the number of nodes is not within the default node range, the number of Task nodes is automatically increased or decreased to the default node range.
  - Time ranges of resource plans cannot be overlapped. The overlapped time range indicates that two effective resource plans exist at a time point. For example, if resource plan 1 takes effect from **08:00** to **10:00** and resource plan 2 takes effect from **09:00** to **11:00**, the time range between **09:00** to **10:00** is overlapped.
  - The time range of a resource plan must be on the same day. For example, if you want to configure a resource plan from **23:00** to **01:00** (the next day), configure two resource plans whose time ranges are **23:00-00:00** and **00:00-01:00**, respectively.
- **Automation scripts**  
When adding an automation script, you can set related parameters by referring to [Table 7-13](#).

**Table 7-13** Configuration items of an automation script

| Configuration Item | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name               | <p>Automation script name.</p> <p>The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and must not start with a space.</p> <p>The value can contain 1 to 64 characters.</p> <p><b>NOTE</b></p> <p>A name must be unique in the same cluster. You can set the same name for different clusters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Script Path        | <p>Script path. The value can be an OBS file system path or a local VM path.</p> <ul style="list-style-type: none"><li>• An OBS file system path must start with <b>obs://</b> and end with <b>.sh</b>, for example, <b>obs://mrs-samples/xxx.sh</b>.</li><li>• A local VM path must start with a slash (/) and end with <b>.sh</b>. For example, the path of the example script for installing the Zepelin is <b>/opt/bootstrap/zepelin/zepelin_install.sh</b>.</li></ul>                                                                                                                                                                                                                                                                                                                                                   |
| Execution Node     | <p>Select a type of the node where an automation script is executed.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• If you select <b>Master</b> nodes, you can choose whether to run the script only on the active Master nodes by enabling or disabling the <b>Active Master</b> switch.</li><li>• If you enable it, the script runs only on the active Master nodes. If you disable it, the script runs on all Master nodes. This switch is disabled by default.</li></ul>                                                                                                                                                                                                                                                                                                                                 |
| Parameter          | <p>Automation script parameter. The following predefined variables can be imported to obtain auto scaling information:</p> <ul style="list-style-type: none"><li>• <b>`\${mrs_scale_node_num}`</b>: Number of auto scaling nodes. The value is always positive.</li><li>• <b>`\${mrs_scale_type}`</b>: Scale-out/in type. The value can be <b>scale_out</b> or <b>scale_in</b>.</li><li>• <b>`\${mrs_scale_node_hostnames}`</b>: Host names of the auto scaling nodes. Use commas (,) to separate multiple host names.</li><li>• <b>`\${mrs_scale_node_ips}`</b>: IP address of the auto scaling nodes. Use commas (,) to separate multiple IP addresses.</li><li>• <b>`\${mrs_scale_rule_name}`</b>: Name of the triggered auto scaling rule. For a resource plan, this parameter is set to <b>resource_plan</b>.</li></ul> |

| Configuration Item  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Executed            | <p>Time for executing an automation script. The following four options are supported: <b>Before scale-out</b>, <b>After scale-out</b>, <b>Before scale-in</b>, and <b>After scale-in</b>.</p> <p><b>NOTE</b><br/>Assume that the execution nodes include Task nodes.</p> <ul style="list-style-type: none"><li>• The automation script executed before scale-out cannot run on the Task nodes to be added.</li><li>• The automation script executed after scale-out can run on the added Task nodes.</li><li>• The automation script executed before scale-in can run on Task nodes to be deleted.</li><li>• The automation script executed after scale-in cannot run on the deleted Task nodes.</li></ul> |
| Action upon Failure | <p>Whether to continue to execute subsequent scripts and scale-out/in after the script fails to be executed.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• You are advised to set this parameter to <b>Continue</b> in the commissioning phase so that the cluster can continue the scale-out/in operation no matter whether the script is executed successfully.</li><li>• If the script fails to be executed, view the log in <code>/var/log/Bootstrap</code> on the cluster VM.</li><li>• The scale-in operation cannot be rolled back. Therefore, the <b>Action upon Failure</b> can only be set to <b>Continue</b> after scale-in.</li></ul>                                         |

 **NOTE**

The automation script is triggered only during auto scaling. It is not triggered when the cluster node is manually scaled out or in.

## 7.5.6.2 Adding an Auto Scaling Policy for MRS Task Nodes

### Configuring Auto Scaling During Cluster Creation

When you create a cluster, you can configure the auto scaling function in advanced configuration parameters.

 **NOTE**

Auto scaling policies can be configured during cluster creation only for analysis, streaming, and hybrid clusters.

**Step 1** Log in to the MRS management console.

**Step 2** When you buy a cluster containing task nodes, configure the cluster software and hardware information by referring to [Manually Buying an MRS Cluster](#). Then, on the **Set Advanced Options** page, enable **Analysis Task** and configure or modify auto scaling rules and resource plans.

**Figure 7-14** Configuring auto scaling rules when creating a cluster

**NOTE**

You can configure the auto scaling rules by referring to the following scenarios:

- [Scenario 1: Using Auto Scaling Rules Alone](#)
- [Scenario 2: Using Resource Plans Alone](#)
- [Scenario 3: Using Auto Scaling Rules and Resource Plans Together](#)

----End

## Adding an Auto Scaling Policy for a Cluster

After a cluster is created, you can configure rules for the task node group in a cluster by node group or resource pool.

The node group policy and resource pool policy are mutually exclusive. You can configure either of them as needed.

MRS 3.1.5 or later supports the specified resource pool policy.

| Item                                   | By Node Group                                   | By Resource Pool                                                                |
|----------------------------------------|-------------------------------------------------|---------------------------------------------------------------------------------|
| Auto scaling object                    | All nodes in the task node group                | Task nodes in the resource pool specified by an auto scaling policy             |
| Resource pool ownership of added nodes | Default resource pool                           | Resource pool specified by the auto scaling policy                              |
| Scale-in object                        | Random scale-in of nodes in the task node group | Random scale-in of nodes in a resource pool specified by an auto scaling policy |



### Prerequisites

- A task node group has been configured by referring to [Adding a Task Node](#).
- A resource pool has been added by referring to [Adding an MRS Tenant Resource Pool](#) if you plan to configure auto scaling policies by resource pool.

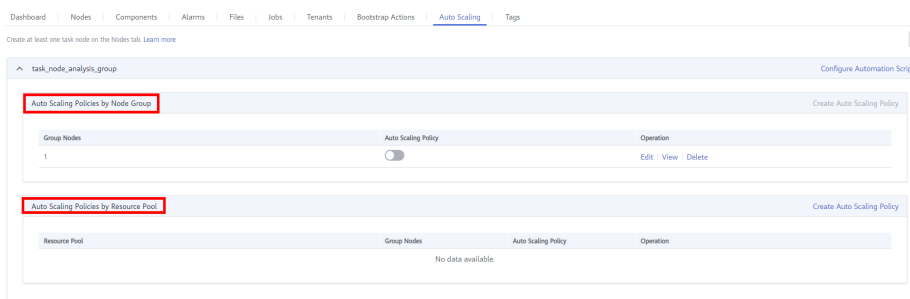
### Procedure

**Step 1** Log in to the MRS management console.

**Step 2** Choose **Active Clusters**, select a running cluster, and click its name to go its details page.

**Step 3** On the page that is displayed, click the **Auto Scaling** tab.

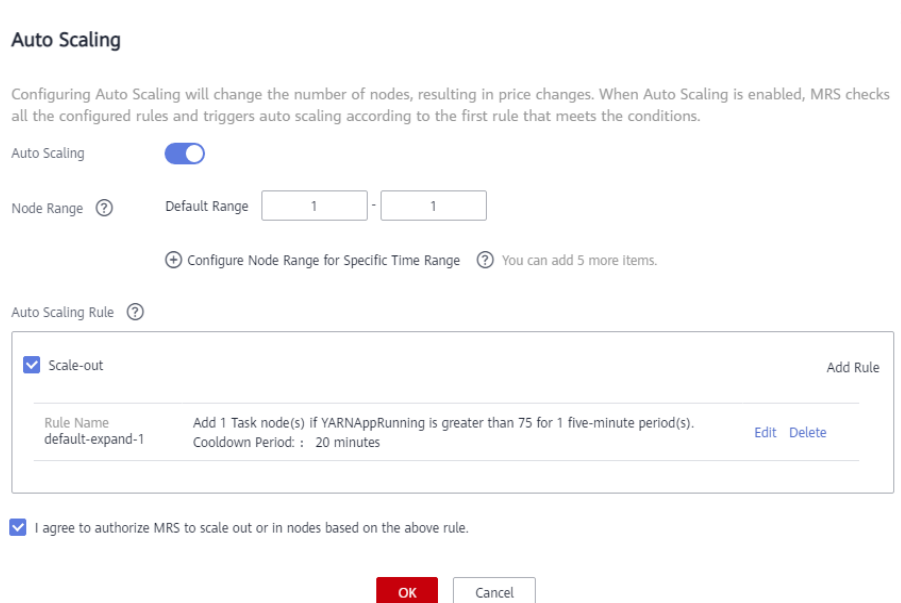
You can configure policies by resource pool or node group as needed.



### NOTE

- Auto scaling policies of different node groups are mutually exclusive. That is, you can enable auto scaling policies only for one node group.
- An auto scaling rule adjusts the number of nodes, but also affects the actual price. Exercise caution when adding an auto scaling rule.

**Step 4** Click **Create Auto Scaling Policy** to create an auto scaling policy.



 NOTE

You can configure the auto scaling rules by referring to the following scenarios:

- [Scenario 1: Using Auto Scaling Rules Alone](#)
- [Scenario 2: Using Resource Plans Alone](#)
- [Scenario 3: Using Auto Scaling Rules and Resource Plans Together](#)

----End


## Scenario 1: Using Auto Scaling Rules Alone

Scenario where only auto scaling rules are configured: The number of nodes needs to be dynamically adjusted based on the YARN resource usage. When the available YARN memory is less than 20%, five nodes need to be added. When the available YARN memory is greater than 70%, five nodes need to be reduced. The number of nodes in a task node group ranges from 1 to 10.

**Step 1** Go to the **Auto Scaling** page to configure auto scaling rules.

- Configure the **Default Range** parameter.  
Enter a task node range, in which auto scaling is performed. This constraint applies to all scale-in and scale-out rules. The maximum value range allowed is 0 to 500.  
The value range in this example is 1 to 10.
- Configure an auto scaling rule.  
To enable **Auto Scaling**, you must configure a scale-out or scale-in rule.
  - Select **Scale-Out** or **Scale-In**.
  - Click **Add Rule**.

**Figure 7-15** Adding a rule



**Add Rule**

Rule Name:

If:

Last For:  five-minute periods

Add:  nodes

Cooldown Period:  minutes

- Configure **Rule Name**, **If**, **Last For**, **Add**, and **Cooldown Period**. For details about auto scaling metrics, see [Automatic Scaling of Task Nodes in an MRS Cluster](#).

- d. Click **OK**.

You can view, edit, or delete the rules you configured in the **Scale-out** or **Scale-in** area on the **Auto Scaling** page. You can click **Add Rule** to configure multiple rules.

**Step 2** Click **OK**.

 **NOTE**

If you want to configure an auto scaling rule for an existing cluster, select **I agree to authorize MRS to scale out or in nodes based on the above rule**.

----End

## Scenario 2: Using Resource Plans Alone

If the data volume changes regularly every day and you want to scale out or in a cluster before the data volume changes, you can create resource plans to adjust the number of Task nodes as planned in the specified time range.

For example: A real-time processing service sees a sharp increase in data volume from 7:00 to 13:00 on Monday, Tuesday, and Saturday. Assume that an MRS streaming cluster is used to process the service data. Five task nodes are required from 7:00 to 13:00 on Monday, Tuesday, and Saturday, while only two are required at other time.

**Step 1** Go to the **Auto Scaling** page to configure a resource plan.

 **NOTE**

A resource plan can be configured to adjust the number of nodes, which affects the actual price. Exercise caution when performing this operation.

**Step 2** For example, the **Default Range** of node quantity is set to **2-2**, indicating that the number of task nodes is fixed to 2 except the time range specified in the resource plan.

**Step 3** Click **Configure Node Range for Specific Time Range** under **Default Range** or **Add Resource Plan**.

**Step 4** Configure **Effective On**, **Time Range**, and **Node Range**.

For example, set **Effective On** to **Monday, Tuesday, and Saturday**, **Time Range** to **07:00-13:00**, and **Node Range** to **5-5**. This indicates that the number of task nodes is fixed at 5 from 07:00 to 13:00.

Click **Configure Node Range for Specific Time Range** to configure multiple resource plans.

 NOTE

- **Effective On** is set to **Daily** by default. You can also select one or multiple days from Monday to Sunday.
- If you do not set **Node Range**, its default value will be used.
- If you set both **Node Range** and **Time Range**, the node range you set will be used during the time range you set, and the default node range will be used beyond the time range you set. If the time is not within the configured time range, the default range is used.

----End

### Scenario 3: Using Auto Scaling Rules and Resource Plans Together

If the data volume is not stable and the expected fluctuation may occur, the fixed Task node range cannot guarantee that the requirements in some service scenarios are met. In this case, it is necessary to adjust the number of Task nodes based on the real-time loads and resource plans.

For example: A real-time processing service sees an unstable increase in data volume from 7:00 to 13:00 on Monday, Tuesday, and Saturday. For example, 5 to 8 task nodes are required from 7:00 to 13:00 on Monday, Tuesday, and Saturday, and 2 to 4 are required beyond this period. You can set an auto scaling rule based on a resource plan. When the data volume exceeds the expected value, the number of Task nodes changes with resource loads, without exceeding the node range specified in the resource plan. When a resource plan is triggered, the number of nodes is adjusted within the specified node range with minimum affect. That is, increase nodes to the upper limit and decrease nodes to the lower limit.

**Step 1** Go to the **Auto Scaling** page to configure auto scaling rules.

An auto scaling rule adjusts the number of nodes, but also affects the actual price. Exercise caution when adding an auto scaling rule.

- **Default Range**  
Enter a task node range, in which auto scaling is performed. This constraint applies to all scale-in and scale-out rules.  
For example, this parameter is set to **2-4** in this scenario.
- **Auto Scaling**  
To enable **Auto Scaling**, you must configure a scale-out or scale-in rule.
  - a. Select **Scale-Out** or **Scale-In**.
  - b. Click **Add Rule**. The **Add Rule** page is displayed.

Figure 7-16 Adding a rule

## Add Rule

|                                                                         |                                                                                                                       |
|-------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Rule Name                                                               | <input type="text" value="default-expand-2"/>                                                                         |
| If                                                                      | <input type="text" value="YARNAppRunning"/> <input type="text" value="Greater than"/> <input type="text" value="75"/> |
| Last For                                                                | <input type="text" value="1"/> five-minute periods                                                                    |
| Add                                                                     | <input type="text" value="1"/> nodes                                                                                  |
| Cooldown Period                                                         | <input type="text" value="20"/> minutes                                                                               |
| <input type="button" value="OK"/> <input type="button" value="Cancel"/> |                                                                                                                       |

- c. Configure the **Rule Name**, **If**, **Last for**, **Add**, and **Cooldown Period** parameters.
- d. Click **OK**.

You can view, edit, or delete the rules you configured in the **Scale-out** or **Scale-in** area on the **Auto Scaling** page.

**Step 2** Configure a resource plan.

1. Click **Configure Node Range for Specific Time Range** under **Default Range** or **Add Resource Plan**.
2. Configure **Effective On**, **Time Range**, and **Node Range**.

For example, set **Effective On** to **Monday, Tuesday, and Saturday**, **Time Range** to **07:00-13:00**, and **Node Range** to **5-8**.

Click **Configure Node Range for Specific Time Range** or **Add Resource Plan** to configure multiple resource plans.

**NOTE**

- **Effective On** is set to **Daily** by default. You can also select one or multiple days from Monday to Sunday.
- If you do not set **Node Range**, its default value will be used.
- If you set both **Node Range** and **Time Range**, the node range you set will be used during the time range you set, and the default node range will be used beyond the time range you set. If the time is not within the configured time range, the default range is used.

----End

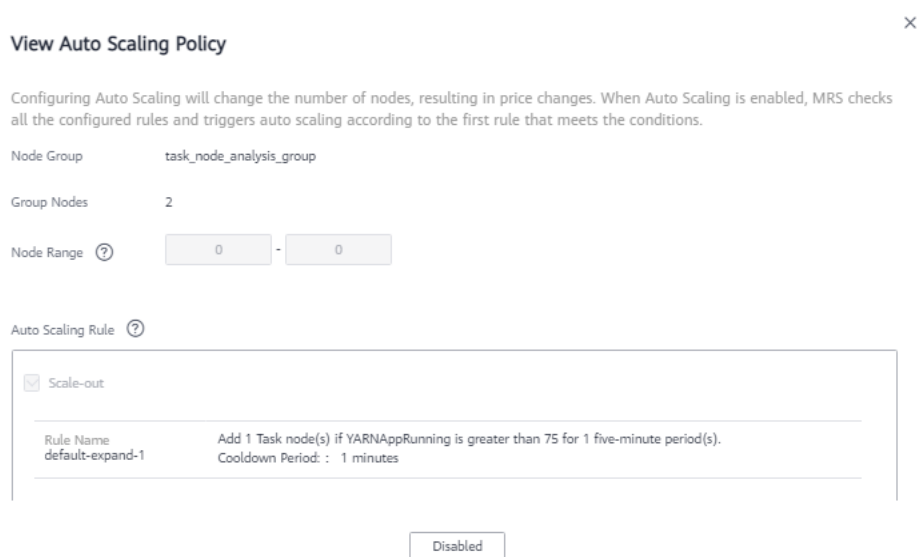
### 7.5.6.3 Managing MRS Cluster Auto Scaling Policies

You can view, modify, delete, enable, and disable an auto scaling policy on the MRS console after it is created.

## Viewing an Auto Scaling Policy

- Step 1** Log in to the MRS management console.
- Step 2** Choose **Active Clusters**, select a running cluster, and click its name. to access its details page.
- Step 3** Click the **Auto Scaling** tab.
- Step 4** Click **View** on the right of the target auto scaling policy to view it.

**Figure 7-17** Viewing an auto scaling policy



----End

## Modifying an Auto Scaling Policy

- Step 1** Log in to the MRS management console.
- Step 2** Choose **Active Clusters**, select a running cluster, and click its name. to access its details page.
- Step 3** Click the **Auto Scaling** tab.
- Step 4** Click **Edit** on the right of the target auto scaling policy.

**Figure 7-18** Editing an auto scaling policy

**Edit Auto Scaling Policy** ✕

Configuring Auto Scaling will change the number of nodes, resulting in price changes. When Auto Scaling is enabled, MRS checks all the configured rules and triggers auto scaling according to the first rule that meets the conditions.

Node Group: task\_node\_analysis\_group

Group Nodes: 2

Node Range:  -

[Configure Node Range for Specific Time Range](#) [You can add 5 more items.](#)

Auto Scaling Rule [?](#)

|                                                                                                                                                                                                                                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <input checked="" type="checkbox"/> Scale-out <span style="float: right;">Add Rule</span>                                                                                                                                          |
| Rule Name: default-expand-1<br>Add 1 Task node(s) if YARNAppRunning is greater than 75 for 1 five-minute period(s).<br>Cooldown Period: : 1 minutes <span style="float: right;"><a href="#">Edit</a> <a href="#">Delete</a></span> |

I agree to authorize MRS to scale out or in nodes based on the above rule.

----End

## Deleting an Auto Scaling Policy

- Step 1** Log in to the MRS management console.
- Step 2** Choose **Active Clusters**, select a running cluster, and click its name. to access its details page.
- Step 3** Click the **Auto Scaling** tab.
- Step 4** Click **Delete** on the right of an existing auto scaling policy. In the displayed dialog box, click **OK**.

----End

## Enabling or Disabling an Auto Scaling Policy

- Step 1** Log in to the MRS management console.
- Step 2** Choose **Active Clusters**, select a running cluster, and click its name. to access its details page.
- Step 3** Click the **Auto Scaling** tab.
- Step 4** Toggle **Auto Scaling Policy** on or off to enable or disable an auto scaling policy.

----End

# 7.6 MRS Cluster Data Backup and Restoration

## 7.6.1 Backing Up and Restoring MRS Cluster Data

### Overview

Manager can back up system and user data by components. The system can back up Manager data, component metadata, and service data.

For MRS 3.x and later, data can be backed up to local disks (LocalDir), local HDFS (LocalHDFS), remote HDFS (RemoteHDFS), NAS (NFS/CIFS), Object Storage Service (OBS), and SFTP server (SFTP). For details, see [Backing Up MRS Cluster Component Data](#).

#### NOTE

Only MRS 3.1.0 or later supports data backup to OBS.

Backup and restoration tasks are performed in the following scenarios:

- Routine backup is performed to ensure the data security of the system and components.
- If the system is faulty, the data backup can be used to recover the system.
- If the active cluster is completely faulty, a mirrored cluster identical to the active cluster needs to be created. You can use the backup data to restore the active cluster.

**Table 7-14** Metadata (MRS 2.x and earlier versions)

| Backup Type | Backup Content                                                                                                          |
|-------------|-------------------------------------------------------------------------------------------------------------------------|
| OMS         | Database data (excluding alarm data) and configuration data in the cluster management system to be backed up by default |
| LdapServer  | User information, including the username, password, key, password policy, and user group information                    |
| DBService   | Metadata of the components (Hive) managed by DBService                                                                  |
| NameNode    | HDFS metadata                                                                                                           |



**Table 7-15** Manager configuration data (MRS 3.x and later)

| Backup Type | Backup Content                                                                                          | Backup Directory Type                                                                                                                                                 |
|-------------|---------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OMS         | Database data (excluding alarm data) and configuration data in the cluster management system by default | <ul style="list-style-type: none"> <li>• LocalDir</li> <li>• LocalHDFS</li> <li>• RemoteHDFS</li> <li>• NFS</li> <li>• CIFS</li> <li>• SFTP</li> <li>• OBS</li> </ul> |

**Table 7-16** Component metadata or other data (MRS 3.x and later)

| Backup Type                                           | Backup Content                                                                                                                                                                             | Backup Directory Type                                                                                                                                                 |
|-------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DBService                                             | Metadata of the components (including Loader, Hive, Spark, Oozie, CDL, and Hue) managed by DBService.                                                                                      | <ul style="list-style-type: none"> <li>• LocalDir</li> <li>• LocalHDFS</li> <li>• RemoteHDFS</li> <li>• NFS</li> <li>• CIFS</li> <li>• SFTP</li> <li>• OBS</li> </ul> |
| Flink<br>(Applicable to MRS 3.2.0 and later versions) | Flink metadata.                                                                                                                                                                            | <ul style="list-style-type: none"> <li>• LocalDir</li> <li>• LocalHDFS</li> <li>• RemoteHDFS</li> </ul>                                                               |
| Kafka                                                 | Kafka metadata.                                                                                                                                                                            | <ul style="list-style-type: none"> <li>• LocalDir</li> <li>• LocalHDFS</li> <li>• RemoteHDFS</li> <li>• NFS</li> <li>• CIFS</li> <li>• OBS</li> </ul>                 |
| NameNode                                              | HDFS metadata. After multiple NameServices are added, backup and restoration are supported for all of them and the operations are consistent with those of the default hacluster instance. | <ul style="list-style-type: none"> <li>• LocalDir</li> <li>• RemoteHDFS</li> <li>• NFS</li> <li>• CIFS</li> <li>• SFTP</li> <li>• OBS</li> </ul>                      |

| Backup Type | Backup Content                                                | Backup Directory Type                                                                                                             |
|-------------|---------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| Yarn        | Information about the Yarn service resource pool.             |                                                                                                                                   |
| HBase       | <b>tableinfo</b> files and data files of HBase system tables. |                                                                                                                                   |
| IoTDB       | IoTDB metadata.                                               | <ul style="list-style-type: none"> <li>• LocalDir</li> <li>• NFS</li> <li>• RemoteHDFS</li> <li>• CIFS</li> <li>• SFTP</li> </ul> |
| ClickHouse  | ClickHouse metadata.                                          | <ul style="list-style-type: none"> <li>• LocalDir</li> <li>• RemoteHDFS</li> </ul>                                                |

**Table 7-17** Service data of specific components (MRS 3.x and later)

| Backup Type | Backup Content                                                                                                  | Backup Directory Type                                                                   |
|-------------|-----------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| HBase       | Table-level user data.                                                                                          | <ul style="list-style-type: none"> <li>• RemoteHDFS</li> </ul>                          |
| HDFS        | Directories or files of user services.<br><b>NOTE</b><br>Encrypted directories cannot be backed up or restored. | <ul style="list-style-type: none"> <li>• NFS</li> <li>• CIFS</li> <li>• SFTP</li> </ul> |
| Hive        | Table-level user data.                                                                                          |                                                                                         |
| IoTDB       | IoTDB service data.                                                                                             | RemoteHDFS                                                                              |
| ClickHouse  | Table-level user data.                                                                                          | RemoteHDFS                                                                              |

Note that some components in MRS 3.x and later versions do not provide data backup or restoration:

- Kafka supports replicas and allows multiple replicas to be specified when a topic is created.
- CDL data is stored in DBService and Kafka. A system administrator can create DBService and Kafka backup tasks to back up data.
- MapReduce and Yarn data is stored in HDFS. Therefore, they rely on the backup and restoration provided by HDFS.
- Backup and restoration of service data in ZooKeeper are performed by their own upper-layer components.

## MRS Cluster Data Backup and Restoration Principles

### Task

Before backup or restoration, you need to create a backup or restoration task and set task parameters, such as the task name, backup data source, and type of the directory for storing backup files. Then you can execute the tasks to back up or restore data. When Manager is used to restore the data of HDFS, HBase (MRS 3.x and later), Hive, and NameNode, the cluster cannot be accessed.

Each backup task can back up data of different data sources and generate an independent backup file for each data source. All the backup files generated in a backup task form a backup file set, which can be used in restoration tasks. Backup data can be stored on Linux local disks, local cluster HDFS, and standby cluster HDFS.

- For MRS 3.x and later versions, backup tasks support full backup and incremental backup policies. Cloud data backup tasks do not support incremental backup. If the backup directory type is NFS or CIFS, incremental backup is not recommended. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.
- For MRS 2.x and earlier versions, the backup task provides the full backup or incremental backup policies. HDFS and Hive backup tasks support the incremental backup policy, while OMS, LdapServer, DBService, and NameNode backup tasks support only the full backup policy.

### NOTE

Task execution rules:

- If a task is being executed, the task cannot be executed repeatedly and other tasks cannot be started at the same time.
- The interval at which a periodic task is automatically executed must be greater than 120s. Otherwise, the task is postponed and will be executed in the next period. Manual tasks can be executed at any interval.
- When a periodic task is to be automatically executed, the current time cannot be 120s later than the task start time. Otherwise, the task is postponed and executed in the next period.
- When a periodic task is locked, it cannot be automatically executed and needs to be manually unlocked.
- The **LocalBackup** partition on the active management node must have at least 20 GB of free space to start the backup tasks for OMS, LdapServer (MRS 2.x or earlier), DBService, Kafka (MRS 3.x or later), and NameNode.
- When you are planning backup and restoration tasks, select the data to be backed up or restored strictly based on the service logic, data store structure, and database or table association.
  - For MRS 2.x and earlier versions, the system creates a default periodic backup task **default** whose execution interval is 24 hours to perform full backup of OMS, LdapServer, DBService, and NameNode data to the Linux local disk.
  - For MRS 3.x and later versions, the system creates periodic backup tasks **default-oms** and **default-cluster ID** at an interval of one hour by default. OMS metadata and cluster metadata, such as DBService and NameNode, can be fully backed up to local disks.

### Snapshot (MRS 3.x and later versions)

The system uses the snapshot technology to quickly back up data. Snapshots include HBase and HDFS snapshots.

- **HBase snapshots**

An HBase snapshot is a backup file of HBase tables at a specified time point. This backup file does not replicate service data or affect the RegionServer. The HBase snapshot replicates table metadata, including table descriptor, region info, and HFile reference information. The metadata can be used to restore data before the snapshot creation time.

- **HDFS snapshots**

An HDFS snapshot is a read-only backup of HDFS at a specified time point. The snapshot is used in data backup, misoperation protection, and disaster recovery scenarios.

The snapshot function can be enabled for any HDFS directory to create the related snapshot file. Before creating a snapshot for a directory, the system automatically enables the snapshot function for the directory. Creating a snapshot does not affect any HDFS operation. A maximum of 65,536 snapshots can be created for each HDFS directory.

When a snapshot is being created for an HDFS directory, the directory cannot be deleted or modified before the snapshot is created. Snapshots cannot be created for the upper-layer directories or subdirectories of the directory.

### **DistCp (MRS 3.x and later versions)**

Distributed copy (DistCp) is a tool used to replicate a large amount of data in HDFS in a cluster or between the HDFSs of different clusters. In a backup or restoration task of HBase, HDFS, or Hive, if you back up the data to HDFS of the standby cluster, the system invokes DistCp to perform the operation. Install the MRS software of the same version for the active and standby clusters and install the cluster.

DistCp uses MapReduce to implement data distribution, troubleshooting, restoration, and report. DistCp specifies different Map jobs for various source files and directories in the specified list. Each Map job copies the data in the partition that corresponds to the specified file in the list.

If you use DistCp to replicate data between HDFSs of two clusters, configure the cross-cluster mutual trust (mutual trust does not need to be configured for clusters managed by the same FusionInsight Manager) and cross-cluster replication for both clusters. When backing up the cluster data to HDFS in another cluster, you need to install the Yarn component. Otherwise, the backup fails.

### **Local Fast Restoration (MRS 3.x and later versions)**

After using DistCp to back up the HBase, HDFS, and Hive data of the local cluster to the HDFS of the standby cluster, the HDFS of the local cluster retains the backup data snapshots. You can create local rapid restoration tasks to restore data by using the snapshot files in the HDFS of the local cluster.

### **NAS (MRS 3.x and later versions)**

Network Attached Storage (NAS) is a dedicated data storage server which includes the storage components and embedded system software. It provides the cross-platform file sharing function. By using NFS (supporting NFSv3 and NFSv4) and CIFS (supporting SMBv2 and SMBv3), you can connect the service plane of MRS to the NAS server to back up data to the NAS or restore data from the NAS.

 NOTE

- Before data is backed up to the NAS, the system automatically mounts the NAS shared address to a local partition of the backup task execution node. After the backup is complete, the system unmounts the NAS shared partition from the backup task execution node.
- To prevent backup and restoration failures, do not access the shared address where the NAS server has been mounted to, for example, `/srv/BigData/LocalBackup/nas`, during data backup and restoration.
- When service data is backed up to the NAS, DistCp is used.

## Specifications of MRS Cluster Data Backup and Restoration

**Table 7-18** Specifications of the backup and restoration feature

| Item                                                    | Specification |
|---------------------------------------------------------|---------------|
| Maximum number of backup or restoration tasks           | 100           |
| Number of concurrent tasks in a cluster                 | 1             |
| Maximum number of waiting tasks                         | 199           |
| Maximum size (GB) of backup files on a Linux local disk | 600           |

 NOTE

If service data is stored in the ZooKeeper upper-layer components in MRS 3.x and later versions, ensure that the number of znodes in a single backup or restoration task is not too large. Otherwise, the task will fail, and the ZooKeeper service performance will be affected. To check the number of znodes in a single backup or restoration task, perform the following operations:

- Ensure that the number of znodes in a single backup or restoration task is smaller than the upper limit of OS file handles. Specifically:
  1. To check the upper limit at the system level, run the `cat /proc/sys/fs/file-max` command.
  2. To check the upper limit at the user level, run the `ulimit -n` command.
- If the number of znodes in the parent directory exceeds the upper limit, back up and restore data in its sub-directories in batches. To check the number of znodes using ZooKeeper client scripts, perform the following operations:
  1. On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services > ZooKeeper > Instance**, and view the management IP address of each ZooKeeper role.
  2. Log in to the node where the client is installed, configure environment variables, authenticate the user (skip this step for clusters with Kerberos authentication disabled), and run the following command:

```
zkCli.sh -server ip:port
```

*ip* can be any management IP address. The default value of *port* is 2181.

3. If the following information is displayed, login to the ZooKeeper server is successful:  

```
WatchedEvent state:SyncConnected type:None path:null
[zk: ip:port(CONNECIED) 0]
```
4. Run the `getusage` command to check the number of znodes in the directory to be backed up.

```
getusage /hbase/region
```

In the command output, **Node count=xxxxxx** indicates the number of znodes stored in the **region** directory.

**Table 7-19** Specifications of the **default** task (MRS 2.x and earlier)

| Item                            | OMS                                                                                    | LdapServer | DBService | NameNode |
|---------------------------------|----------------------------------------------------------------------------------------|------------|-----------|----------|
| Backup period                   | 1 hour                                                                                 |            |           |          |
| Maximum number of backup copies | 2                                                                                      |            |           |          |
| Maximum size of a backup file   | 10 MB                                                                                  | 20 MB      | 100 MB    | 1.5 GB   |
| Maximum size of disk space used | 20 MB                                                                                  | 40 MB      | 200 MB    | 3 GB     |
| Save path of backup data        | <i>Data save path</i> / <b>LocalBackup/</b> of the active and standby management nodes |            |           |          |

**Table 7-20** Specifications of the **default** task (MRS 3.x and later)

| Item                            | OMS                                                                              | HBase   | Kafka  | DBService | NameNode                     |
|---------------------------------|----------------------------------------------------------------------------------|---------|--------|-----------|------------------------------|
| Backup period                   | 1 hour                                                                           |         |        |           |                              |
| Maximum number of backups       | 168 (7-day historical data)                                                      |         |        |           | 24 (one-day historical data) |
| Maximum size of a backup file   | 10 MB                                                                            | 10 MB   | 512 MB | 100 MB    | 20 GB                        |
| Maximum size of disk space used | 1.64 GB                                                                          | 1.64 GB | 84 GB  | 16.41 GB  | 480 GB                       |
| Storage path of backup data     | <i>Data storage path/LocalBackup/</i> of the active and standby management nodes |         |        |           |                              |

**NOTE**

- The backup data of the default backup task must be periodically transferred and saved outside the cluster based on the enterprise O&M requirements.
- In MRS 3.x and later, administrators can create DistCp backup tasks to save OMS, DBService, and NameNode data to external clusters.
- The execution time of a cluster data backup task can be calculated using the following formula: Task execution time = Volume of data to be backed up/Network bandwidth between the cluster and the backup device. In practice, you are advised to multiply the calculated duration by 1.5 to get the reference value of the task execution time.
- Executing a data backup task affects the maximum I/O performance of the cluster. Therefore, you are advised to execute a backup task during off-peak hours.

## 7.6.2 Enabling MRS Inter-Cluster Replication

DistCp is used to replicate the data stored in HDFS from a cluster to another cluster. To use DistCp, you must first enable inter-cluster replication on both nodes of the cluster where you want to copy data.

Administrators can modify parameters on Manager to enable inter-cluster replication. They then create a backup task to copy data to the remote HDFS.

### Impact on the System

Yarn needs to be restarted to enable the cross-cluster replication function and cannot be accessed during restart.

### Prerequisites

- The **hadoop.rpc.protection** parameter of HDFS in the two clusters for data replication must use the same data transmission mode. The default value is

**privacy**, indicating encrypted transmission. The value **authentication** indicates that transmission is not encrypted.

- For clusters with Kerberos authentication enabled (security mode), mutual trust between clusters needs to be configured.
- The inbound rules of the two security groups on the peer cluster have been added to the two security groups in each cluster to allow all access requests of all protocols and ports of all ECSs in the security groups.

## Enabling MRS Inter-Cluster Replication

**Step 1** Log in to the Manager of one of the two clusters.

- For MRS 2.x and earlier, choose **Services > Yarn > Service Configuration** and set **Type** to **All**.
- For MRS 3.x and later, choose **Cluster > Services > Yarn > Configurations**, and click **All Configurations**.

**Step 2** In the navigation pane on the left, choose **Yarn > Distcp** and set the following parameters:

- For MRS2.x and earlier versions, set **dfs.namenode.rpc-address.haclusterX.remotenn1** to the service IP address and RPC port of a NameNode instance in the peer cluster, in **dfs.namenode.rpc-address.haclusterX.remotenn2**, enter the service IP address and RPC port number of the other NameNode instance in the peer cluster. For example, enter **10.1.1.1:25000** and **10.1.1.2:25000**.  
**dfs.namenode.rpc-address.haclusterX.remotenn1** and **dfs.namenode.rpc-address.haclusterX.remotenn2** do not distinguish active and standby NameNode instances. The default NameNode RPC port is **25000** and cannot be modified on Manager.
- For MRS 3.x and later versions, modify **dfs.namenode.rpc-address**, set **haclusterX.remotenn1** to the service IP address and RPC port of one NameNode instance of the peer cluster, and set **haclusterX.remotenn2** to the service IP address and RPC port number of the other NameNode instance of the peer cluster. Examples of modified parameter values: **10.1.1.1:8020** and **10.1.1.2:8020**.  
**haclusterX.remotenn1** and **haclusterX.remotenn2** do not distinguish active and standby NameNodes. The default NameNode RPC port is **8020** and cannot be modified on Manager.

### NOTE

If data of the current cluster needs to be backed up to the HDFS of multiple clusters, you can configure the corresponding NameNode RPC addresses to **haclusterX1**, **haclusterX2**, **haclusterX3**, and **haclusterX4**.

**Step 3** Save the configurations and restart YARN.

- For MRS2.x and earlier versions, click **Save Configuration** and select **Restart the affected services or instances**. Click **OK** to restart the YARN service. After the system displays "Operation succeeded", click **Finish**. The YARN service is restarted successfully.
- For MRS 3.x and later versions, click **Save**. In the displayed dialog box, click **OK**. Click **Dashboard**, choose **More > Restart Service**, and enter the password of the user to restart the YARN service.



**Step 4** Log in to Manager of the other cluster and repeat [Step 1](#) to [Step 3](#).

----End

## 7.6.3 Creating an MRS Cluster Data Backup Task

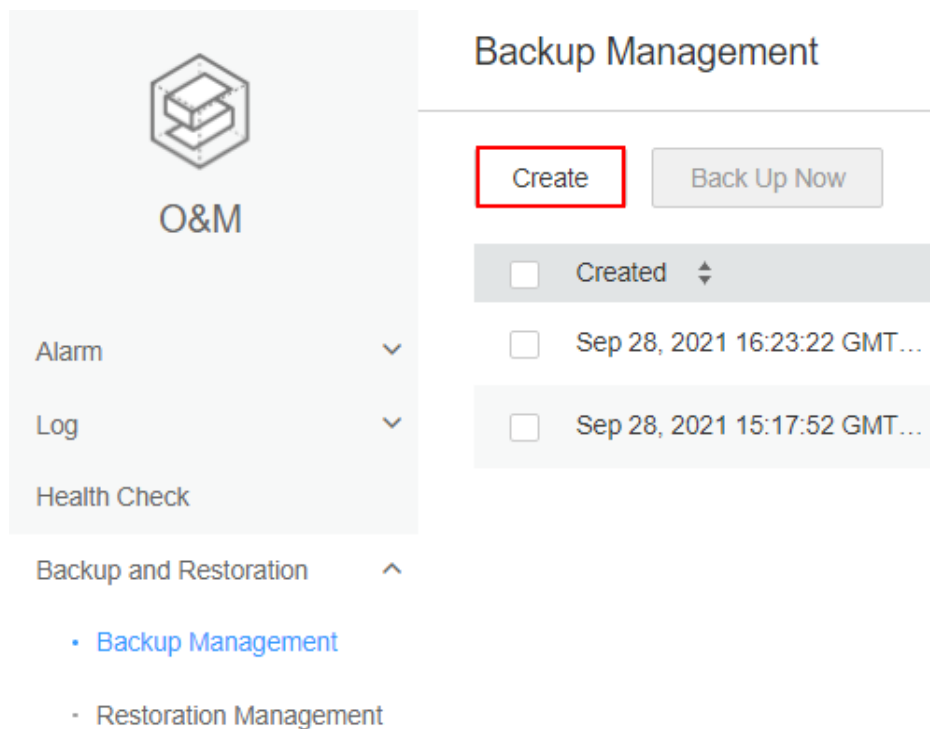
You can create backup tasks on Manager. Executing backup tasks backs up related data.

### Creating a Data Backup Task (MRS 3.x and Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Backup and Restoration > Backup Management**. On the page that is displayed, click **Create**.

**Figure 7-19** Creating a backup task.



**Step 3** Set **Backup Object** to **OMS** or the cluster whose data you want to back up.

**Step 4** Enter a task name in the **Name** text box.

**Step 5** Set **Mode** to **Periodic** or **Manual** as required.

**Table 7-21** Backup types

| Type            | Parameter  | Description                                                                   |
|-----------------|------------|-------------------------------------------------------------------------------|
| Periodic backup | Start Time | Indicates the time when a periodic backup task is started for the first time. |

| Type          | Parameter     | Description                                                                                                                                                                                                                         |
|---------------|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | Period        | Task execution interval. Value options are <b>Hours</b> and <b>Days</b> .                                                                                                                                                           |
|               | Backup Policy | The following policies can be selected: <ul style="list-style-type: none"><li>• Full backup at the first time and subsequent incremental backup</li><li>• Full backup every time</li><li>• Full backup once every n times</li></ul> |
| Manual backup | N/A           | You need to manually execute the task to back up data.                                                                                                                                                                              |

**Step 6** Set required parameters in the **Configuration** area.

- Metadata and service data can be backed up.
- For details about how to back up data of different components, see [MRS Cluster Data Backup and Restoration](#).

**Step 7** Click **OK** to save the configurations.

**Step 8** In the backup task list, you can view the created backup task.

Locate the row that contains the target backup task, choose **More > Back Up Now** in the **Operation** column to execute the task immediately.

----End

## Creating a Metadata Backup Task (MRS 2.x and Earlier)

**Step 1** Create a backup task.

1. On MRS Manager, choose **System > Back Up Data**.
2. Click **Create Backup Task**.

**Step 2** Configure a backup policy.

1. Set **Name** to the name of the backup task.
2. Set **Backup Mode** to the type of the backup task. **Periodic** indicates that the backup task is periodically executed. **Manual** indicates that the backup task is manually executed.

To create a periodic backup task, set the following parameters:

- **Started**: indicates the time when the task is started for the first time.
- **Period**: indicates the task execution interval. The options include **Hours** and **Days**.
- **Backup Policy**: indicates the volume of data to be backed up in each task execution. The options include **Full backup at the first time and incremental backup later**, **Full backup every time**, and **Full backup once every n times**. If you select **Full backup once every n times**, you need to specify the value of *n*.

**Step 3** Select backup sources.

On the **Configuration** page, select the metadata option and set backup parameters. For details, see [Backing Up Manager Data \(MRS 2.x and Earlier\)](#).

**Step 4** Click **OK**.

**Step 5** In the **Operation** column of the created task in the backup task list, click **Back Up Now** if **Backup Mode** is set to **Periodic** or click **Start** if **Backup Mode** is set to **Manual** to execute the backup task.

----End

## 7.6.4 Creating an MRS Cluster Data Restoration Task

You can create a restoration task on Manager. After the restoration task is executed, the specified backup data is restored to the cluster.

### Creating a Data Restoration Task (MRS 3.x and Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Backup and Restoration > Restoration Management**. On the page that is displayed, click **Create**.

**Step 3** Configure **Task Name**.

**Step 4** Set **Recovery Object** to **OMS** or the cluster whose data you want to restore.

**Step 5** Set the required parameters in the **Recovery Configuration** area.

- Metadata and service data can be restored.
- For details about how to restore data of different components, see [MRS Cluster Data Backup and Restoration](#).

**Step 6** Click **OK** to save the configurations.

**Step 7** In the restoration task list, you can view the created restoration tasks.

Locate the row containing the target restoration task, click **Start** in the **Operation** column to execute the restoration task immediately.

----End

### Creating a Metadata Restoration Task (MRS 2.x and Earlier)

**Step 1** On MRS Manager, choose **System > Recovery Management**.

**Step 2** On the page that is displayed, click **Create Restoration Task**.

**Step 3** Set **Task Name** to the name of the restoration task.

**Step 4** In the **Configuration** area, select the component whose metadata is to be restored and set restoration parameters. For details, see [Restoring Manager Data \(MRS2.x and Earlier\)](#).

**Step 5** Click **OK**.

**Step 6** In the restoration task list, you can view the created restoration tasks.

In the restoration task list, locate the row where the created task resides, and click **Start** in the **Operation** column.

----End

## 7.6.5 Backing Up MRS Cluster Component Data

### 7.6.5.1 Backing Up Manager Data (MRS 2.x and Earlier)

#### Scenario

To ensure the security of metadata either on a routine basis or before and after performing critical metadata operations (such as scale-out, scale-in, patch installation, upgrades, and migration), metadata must be backed up. The backup data can be used to recover the system if an exception occurs or if the operation has not achieved the expected result. This minimizes the adverse impact on services. Metadata includes data of OMS, LdapServer, DBService, and NameNode. MRS Manager data to be backed up includes OMS data and LdapServer data.

By default, metadata backup is supported by the **default** task. This section describes how to create a backup task and back up metadata on MRS Manager. Both automatic backup tasks and manual backup tasks are supported.

#### Prerequisites

- A standby cluster for backing up data has been created, and the network is connected. The inbound rules of the two security groups on the peer cluster have been added to the two security groups in each cluster to allow all access requests of all protocols and ports of all ECSs in the security groups.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.

### Backing Up Manager Data

**Step 1** Create a backup task.

1. On MRS Manager, choose **System > Back Up Data**.
2. Click **Create Backup Task**.

**Step 2** Configure a backup policy.

1. Set **Task Name** to the name of the backup task.
2. Set **Backup Mode** to the type of the backup task. **Periodic** indicates that the backup task is periodically executed. **Manual** indicates that the backup task is manually executed.

To create a periodic backup task, set the following parameters:

- **Started**: indicates the time when the task is started for the first time.
- **Period**: indicates the task execution interval. The options include **By hour** and **By day**.

- **Backup Policy:** indicates the volume of data to be backed up in each task execution. The options include **Full backup at the first time and incremental backup later**, **Full backup every time**, and **Full backup once every n times**. If you select **Full backup once every n times**, you need to specify the value of *n*.

**Step 3** Select backup sources.

In the **Configuration** area, select **OMS** and **LdapServer** under **Metadata**.

**Step 4** Set backup parameters.

1. Set **Path Type** of **OMS** and **LdapServer** to a backup directory type.

The following backup directory types are supported:

- **LocalDir:** indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files. By default, the backup files are stored in *Data storage path/LocalBackup/*. If you select **LocalDir**, you need to set the maximum number of copies to specify the number of backup files that can be retained in the backup directory.
- **LocalHDFS:** indicates that the backup files are stored in the HDFS directory of the current cluster. If you select **SFTP**, set the following parameters:
  - **Target Path:** indicates the HDFS directory for storing the backup files. The save path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory.
  - **Max Number of Backup Copies:** indicates the number of backup files that can be retained in the backup directory.
  - **Target Instance Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.

2. Click **OK**.

**Step 5** Execute the backup task.

In the **Operation** column of the created task in the backup task list, click **Back Up Now** if **Backup Mode** is set to **Periodic** or click **Start If Backup Mode** is set to **Manual** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version\_Data source\_Task execution time.tar.gz*.

----End

## 7.6.5.2 Backing Up Manager Data (MRS 3.x and Later Versions)

### Scenario

To ensure data security of FusionInsight Manager routinely or before and after a critical operation (such as capacity expansion and reduction) on FusionInsight

Manager, you need to back up FusionInsight Manager data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up Manager data. Both automatic and manual backup tasks are supported.

## Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

## Backing Up Manager Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Figure 7-20** Creating a backup task.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Set **Backup Object** to **OMS**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.  
**Manual** indicates that the backup task is executed manually.

**Table 7-22** Periodic backup parameters

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Started       | Indicates the time when the task is started for the first time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Period        | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Backup Policy | <ul style="list-style-type: none"> <li>● <b>Full backup at the first time and incremental backup subsequently</b></li> <li>● <b>Full backup every time</b></li> <li>● <b>Full backup once every n times</b></li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>● Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li> <li>● If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li> </ul> |

**Step 6** In **Configuration**, select **OMS**.

**Step 7** Set **Path Type** of **OMS** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

The default storage directory is *Data storage path/LocalBackup/*, for example, */srv/BigData/LocalBackup/*.

If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.

- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

If you select this option, set the following parameters:

- **Target Path**: indicates the HDFS directory for storing the backup files. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as */hbase* or */user/hbase/backup*.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Cluster for Backup**: Enter the cluster name mapping to the backup directory.
- **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as */hbase* or */user/hbase/backup*.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Source Cluster**: Select the cluster of the Yarn queue used by the backup data.



- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the server where the backup data is stored.
  - **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username:** indicates the username for connecting to the server using the SFTP protocol.
  - **Password:** indicates the password for connecting to the server using the SFTP protocol.

- **Server Shared Path:** indicates the backup path on the SFTP server.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **OBS:** indicates that backup files are stored in OBS.  
If you select this option, set the following parameters:
  - **Target Path:** indicates the OBS directory for storing backup data.
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.

 **NOTE**

Only MRS 3.1.0 or later supports data backup to OBS.

**Step 8** Click **OK**.

**Step 9** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version\_Data source\_Task execution time.tar.gz*.

----End

### 7.6.5.3 Backing Up CDL Service Data

#### Scenario

To ensure CDL service data security routinely or before a major operation on CDL (such as upgrade or migration), you need to back up CDL data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

CDL data is stored in DBService and Kafka. You can create DBService and Kafka backup tasks on FusionInsight Manager to back up CDL data. Both automatic and manual backup tasks are supported.

#### Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.

- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

## Backing Up CDL Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.

**Manual** indicates that the backup task is executed manually.

**Table 7-23** Periodic backup parameters

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Started       | Indicates the time when the task is started for the first time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Period        | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Backup Policy | <ul style="list-style-type: none"><li>• <b>Full backup at the first time and incremental backup subsequently</b></li><li>• <b>Full backup every time</b></li><li>• <b>Full backup once every n times</b></li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li><li>• If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li></ul> |

**Step 6** Set **Configuration** to **DBService** and **Kafka**.

**Step 7** Set **Path Type** of **DBService** to a backup directory type. For details about how to set the parameters, see [Step 7](#).

**Step 8** Set **Path Type** of **Kafka** to a backup directory type. For details about how to set the parameters, see [Step 7](#).

**Step 9** Click **OK**.

**Step 10** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version\_Data source\_Task execution time.tar.gz*.

----End

## 7.6.5.4 Backing Up ClickHouse Metadata

### Scenario

To ensure ClickHouse metadata security or before a major operation (such as upgrade or migration), you need to back up ClickHouse metadata. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up ClickHouse metadata. Both automatic and manual backup tasks are supported.

---

**NOTICE**

This function is supported only by MRS 3.1.0 or later.

---

### Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
  - In active/standby mode, ensure that the value of **HADOOP\_RPC\_PROTECTION** of ClickHouse is the same as that of **hadoop.rpc.protection** of HDFS.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.

## Backing Up ClickHouse Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task. **Periodic** indicates that the backup task is periodically executed. **Manual** indicates that the backup task is manually executed.

To create a periodic backup task, set the following parameters:

- **Started:** indicates the time when the task is started for the first time.
- **Period:** indicates the task execution interval. The options include **Hours** and **Days**.
- **Backup Policy:** Only **Full backup every time** is supported.

**Step 6** In **Configuration**, select **ClickHouse** under **Metadata and other data**.

**Step 7** Set **Path Type** of **ClickHouse** to a backup directory type.

The following backup directory types are supported:

- **LocalDir:** indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

The default storage directory is *Data storage path/LocalBackup/*, for example, */srv/BigData/LocalBackup*.

If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.

- **RemoteHDFS:** indicates that the backup files are stored in the HDFS directory of the standby cluster.

This option is available for MRS 3.1.0 and 3.1.2 clusters only after you configure the environment by referring to [Configuring the Environment When the ClickHouse Backup Task's Path is RemoteHDFS](#).

You also need to configure the following parameters for MRS 3.2.0 or later clusters:

- **Destination NameService Name:** indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Destination Active NameNode IP Address:** indicates the service plane IP address of the active NameNode in the standby cluster.
- **Destination Standby NameNode IP Address:** indicates the service plane IP address of the standby NameNode in the standby cluster.

- **Destination NameNode RPC Port:** indicates the value of `dfs.namenode.rpc.port` in the HDFS basic configuration of the destination cluster.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as `/hbase` or `/user/hbase/backup`.

You also need to configure the following parameters for MRS 3.1.0 or 3.1.2 clusters:

- **Destination NameService Name:** indicates the NameService name of the standby cluster, for example, `hacluster`. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, `IPv4` or `IPv6`.
- **Target NameNode IP Address:** indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as `/hbase` or `/user/hbase/backup`.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.

**Step 8** Click **OK**.

**Step 9** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Data source\_Task execution time.tar.gz*.

----End

## Configuring the Environment When the ClickHouse Backup Task's Path is RemoteHDFS

This topic is available for MRS 3.1.0 and 3.1.2 only.

**Step 1** Log in to FusionInsight Manager of the standby cluster.

**Step 2** Choose **Cluster > Services > HDFS**. Click **More** and select **Download Client**. Set **Select Client Type** to **Configuration Files Only**, select **x86\_64** for x86 or **aarch64** for Arm based on the type of the node, and click **OK**.

**Step 3** After the client file package is generated, download the client to the local PC as prompted and decompress the package.

For example, if the client file package is `FusionInsight_Cluster_1_HDFS_Client.tar`, decompress it to obtain


**FusionInsight\_Cluster\_1\_HDFS\_ClientConfig\_ConfigFiles.tar**. Then, decompress this file to the **D:\FusionInsight\_Cluster\_1\_HDFS\_ClientConfig\_ConfigFiles** directory on the local PC. The directory name cannot contain spaces.

- Step 4** Go to the **FusionInsight\_Cluster\_1\_HDFS\_ClientConfig\_ConfigFiles\ client** directory and obtain the **hosts** file.
- Step 5** Go to **FusionInsight\_Cluster\_1\_HDFS\_ClientConfig\_ConfigFiles\HDFS\config** to obtain the **core-site.xml** and **hdfs-site.xml** files.
- Step 6** Log in to FusionInsight Manager of the source cluster.
- Step 7** Choose **Cluster > Services > ClickHouse**, choose **Configurations > All Configurations**, and select **backup** under **ClickHouse(Service)**.

For **remote\_connect\_core\_config\_file**, click **Upload File** and select the **core-site.xml** file prepared in [Step 5](#).

For **remote\_connect\_hdfs\_config\_file**, click **Upload File** and select the **hdfs-site.xml** file prepared in [Step 5](#).

**Figure 7-21** Configuring ClickHouse data backup parameters



- Step 8** Click **Save**, confirm the information, and click **OK** to save the configuration. After saving the configurations, click **Finish**.
- Step 9** Choose **Cluster > Services > ClickHouse**, click **Instances**, and view the instance IP address of **ClickHouseServer**.
- Step 10** Log in to the host nodes of the ClickHouseServer instances as user **root** and check whether the **/etc/hosts** file contains the host information in [Step 4](#). If not, add the host information in [Step 4](#) to the **/etc/hosts** file.

----End

## 7.6.5.5 Backing Up ClickHouse Service Data

### Scenario

To ensure ClickHouse service data security routinely or before a major operation on ClickHouse (such as upgrade or migration), you need to back up ClickHouse service data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up ClickHouse service data. Both automatic and manual backup tasks are supported.

**NOTICE**

This function is supported only by MRS 3.1.0 or later.

## Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
  - The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.
  - In active/standby mode, ensure that the value of **HADOOP\_RPC\_PROTECTION** of ClickHouse is the same as that of **hadoop.rpc.protection** of HDFS.
- You have planned the backup type, period, object, and directory based on service requirements.

## Backing Up ClickHouse Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.

**Manual** indicates that the backup task is executed manually.

**Table 7-24** Periodic backup parameters

| Parameter | Description                                                                               |
|-----------|-------------------------------------------------------------------------------------------|
| Started   | Indicates the time when the task is started for the first time.                           |
| Period    | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> . |



| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                            |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Policy | <ul style="list-style-type: none"><li>● <b>Full backup at the first time and incremental backup subsequently</b></li><li>● <b>Full backup every time</b></li><li>● <b>Full backup once every n times</b></li></ul> <p><b>NOTE</b><br/>Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</p> |

**Step 6** In **Configuration**, select **ClickHouse** under **Service Data**.

**Step 7** Set **Path Type** of **ClickHouse** to a backup directory type.

Currently, only the **RemoteHDFS** type is available.

**RemoteHDFS**: indicates that backup files are stored in HDFS of the standby cluster.

This option is available for MRS 3.1.0 and 3.1.2 clusters only after you configure the environment by referring to [Configuring the Environment When the ClickHouse Backup Task's Path is RemoteHDFS](#).

You also need to configure the following parameters for MRS 3.2.0 and later clusters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Destination Active NameNode IP Address**: indicates the service plane IP address of the active NameNode in the standby cluster.
- **Destination Standby NameNode IP Address**: indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Destination NameNode RPC Port**: indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the destination cluster.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.

You also need to configure the following parameters for MRS 3.1.0 or 3.1.2 clusters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.

- **Target NameNode IP Address:** indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.

**Step 8** Set **Maximum Number of Recovery Points** to the number of snapshots that can be retained in the cluster.

**Step 9** Set **Backup Content** to one or multiple ClickHouse tables to be backed up.

You can select backup data using either of the following methods:

- Adding a backup data file  
Click the name of a database in the navigation pane to show all the tables in the database, and select specified tables.  
MRS 3.2.0 or later:
  - a. Click **Add**.
  - b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.
  - c. Click **OK**.
- MRS 3.2.0 or later: Regular expression filtering
  - a. Click **Query Regular Expression**.
  - b. Enter the logical cluster and database to which the ClickHouse table belongs in the first text box as prompted. The logical cluster and database must match the existing logical cluster and database, for example, **/default\_cluster/database**.
  - c. Enter a regular expression in the second box. Standard regular expressions are supported. For example, to filter all tables that contain the keyword **test** in the database, enter **test.\***.
  - d. Click **Refresh** to view the displayed tables in **Directory Name**.
  - e. Click **Synchronize** to save the result.

#### NOTE

- When entering regular expressions, click **+** or **-** to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.
- Versions earlier than MRS 3.2.0: Regular expression filtering
  - a. Click **Query Regular Expression**.
  - b. Enter the database where the ClickHouse tables are located in the first text box as prompted. The database must be the same as the existing database, for example, **default**.

- c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all tables in the database, enter `([\s\S]*?)`. To get tables whose names consist of letters and digits, for example, `tb 1`, enter `tb\d*`.
- d. Click **Refresh** to view the displayed tables in **Directory Name**.
- e. Click **Synchronize** to save the result.

 **NOTE**

- When entering regular expressions, click **+** or **-** to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.

**Step 10** Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The directory or table to be backed up does not exist.
- The name of the NameService is incorrect.

**Step 11** Click **OK**.

**Step 12** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Data source\_Task creation time*, and the subdirectory is used to save latest data source backup files.

----End

## Configuring the Environment When the ClickHouse Backup Task's Path is RemoteHDFS

This topic is available for MRS 3.1.0 and 3.1.2 only.

**Step 1** Log in to FusionInsight Manager of the standby cluster.

**Step 2** Choose **Cluster > Services > HDFS**. Click **More** and select **Download Client**. Set **Select Client Type** to **Configuration Files Only**, select **x86\_64** for x86 or **aarch64** for Arm based on the type of the node, and click **OK**.

**Step 3** After the client file package is generated, download the client to the local PC as prompted and decompress the package.

For example, if the client file package is **FusionInsight\_Cluster\_1\_HDFS\_Client.tar**, decompress it to obtain **FusionInsight\_Cluster\_1\_HDFS\_ClientConfig\_ConfigFiles.tar**. Then, decompress this file to the **D:\FusionInsight\_Cluster\_1\_HDFS\_ClientConfig\_ConfigFiles** directory on the local PC. The directory name cannot contain spaces.

**Step 4** Go to the **FusionInsight\_Cluster\_1\_HDFS\_ClientConfig\_ConfigFiles\** client directory and obtain the **hosts** file.

**Step 5** Go to **FusionInsight\_Cluster\_1\_HDFS\_ClientConfig\_ConfigFiles\HDFS\config** to obtain the **core-site.xml** and **hdfs-site.xml** files.

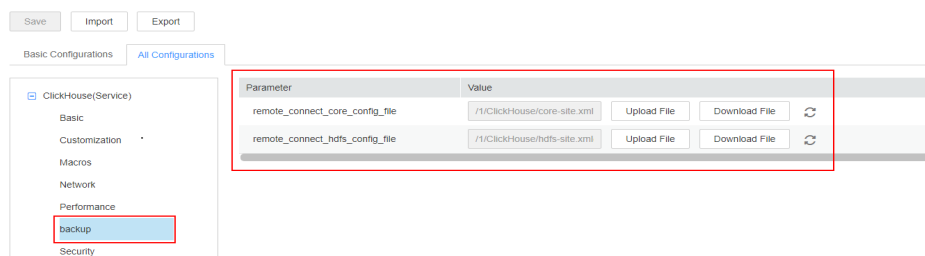
**Step 6** Log in to FusionInsight Manager of the source cluster.

**Step 7** Choose **Cluster > Services > ClickHouse**, choose **Configurations > All Configurations**, and select **backup** under **ClickHouse(Service)**.

For **remote\_connect\_core\_config\_file**, click **Upload File** and select the **core-site.xml** file prepared in **Step 5**.

For **remote\_connect\_hdfs\_config\_file**, click **Upload File** and select the **hdfs-site.xml** file prepared in **Step 5**.

**Figure 7-22** Configuring ClickHouse data backup parameters



**Step 8** Click **Save**, confirm the information, and click **OK** to save the configuration. After saving the configurations, click **Finish**.

**Step 9** Choose **Cluster > Services > ClickHouse**, click **Instances**, and view the instance IP address of **ClickHouseServer**.

**Step 10** Log in to the host nodes of the ClickHouseServer instances as user **root** and check whether the **/etc/hosts** file contains the host information in **Step 4**. If not, add the host information in **Step 4** to the **/etc/hosts** file.

----End

## 7.6.5.6 Backing Up DBService Data

### Scenario

To ensure DBService service data security routinely or before a major operation on DBService (such as upgrade or migration), you need to back up DBService data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up DBService data. Both automatic and manual backup tasks are supported.

### Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager,

mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.

- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

## Backing Up DBService Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.

**Manual** indicates that the backup task is executed manually.

**Table 7-25** Periodic backup parameters

| Parameter | Description                                                                               |
|-----------|-------------------------------------------------------------------------------------------|
| Started   | Indicates the time when the task is started for the first time.                           |
| Period    | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> . |

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Policy | <ul style="list-style-type: none"><li>● <b>Full backup at the first time and incremental backup subsequently</b></li><li>● <b>Full backup every time</b></li><li>● <b>Full backup once every n times</b></li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>● Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li><li>● If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li></ul> |

**Step 6** In **Configuration**, select **DBService**.

**Step 7** Set **Path Type** of **DBService** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

The default storage directory is *Data storage path/LocalBackup/*, for example, */srv/BigData/LocalBackup*.

If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.

- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

If you select this option, set the following parameters:

- **Target Path**: indicates the HDFS directory for storing the backup files. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address:** indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.

If you select this option, set the following parameters:

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address:** indicates the IP address of the server where the backup data is stored.
- **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Server Shared Path:** indicates the backup path on the SFTP server.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **OBS:** indicates that backup files are stored in OBS.

If you select this option, set the following parameters:

- **Target Path:** indicates the OBS directory for storing backup data.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.

 **NOTE**

Only MRS 3.1.0 or later supports data backup to OBS.

**Step 8** Click **OK**.

**Step 9** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version\_Data source\_Task execution time.tar.gz*.

----End

## 7.6.5.7 Backing Up Flink Metadata

### Scenario

To ensure Flink metadata security or before a major operation on Flink (such as upgrade or migration), you need to back up Flink metadata. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up Flink metadata. Both automatic and manual backup tasks are supported.



## Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The HDFS and Yarn services have been installed if data needs to be backed up to HDFS.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.

## Backing Up Flink Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task. **Periodic** indicates that the backup task is periodically executed. **Manual** indicates that the backup task is manually executed.

To create a periodic backup task, set the following parameters:

- **Started**: indicates the time when the task is started for the first time.
- **Period**: indicates the task execution interval. The options include **Hours** and **Days**.
- **Backup Policy**: Only **Full backup every time** is supported.

**Step 6** In **Configuration**, select **Flink** under **Metadata and other data**.

**Step 7** Set **Path Type** of **Flink** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

The default storage directory is *Data storage path/LocalBackup/*, for example, */srv/BigData/LocalBackup*.

If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.

- **LocalHDFS:** indicates that the backup files are stored in the HDFS directory of the current cluster.

If you select this option, set the following parameters:

- **Target Path:** indicates the HDFS directory for storing the backup files. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.

- **RemoteHDFS:** indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name:** indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address:** indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.

**Step 8** Click **OK**.

**Step 9** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Data source\_Task execution time.tar.gz*.

----End

## 7.6.5.8 Backing Up HBase Metadata

### Scenario

To ensure HBase metadata security (including tableinfo files and HFiles) or before a major operation on HBase system tables (such as upgrade or migration), you need to back up HBase metadata to prevent HBase service unavailability caused by HBase system table directory or file damages. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up HBase metadata. Both automatic and manual backup tasks are supported.

### Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- The **fs.defaultFS** parameter settings of HBase are the same as those of Yarn and HDFS.
- If HBase data is stored in the local HDFS, HBase metadata can be backed up to OBS. If HBase data is stored in OBS, data backup is not supported.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

### Backing Up HBase Metadata

- Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.
- Step 2** Click **Create**.
- Step 3** Set **Name** to the name of the backup task.
- Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.  
**Manual** indicates that the backup task is executed manually.

**Table 7-26** Periodic backup parameters

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Started       | Indicates the time when the task is started for the first time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Period        | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Backup Policy | <ul style="list-style-type: none"><li>• <b>Full backup at the first time and incremental backup subsequently</b></li><li>• <b>Full backup every time</b></li><li>• <b>Full backup once every n times</b></li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li><li>• If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li></ul> |

**Step 6** In **Configuration**, select **HBase** under **Metadata and other data**.

**Step 7** Set **Path Type** of **HBase** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.  
The default storage directory is *Data storage path/LocalBackup/*, for example, */srv/BigData/LocalBackup*.  
If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.  
If you select this option, set the following parameters:
  - **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.

- **Target NameNode IP Address:** indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as `/hbase` or `/user/hbase/backup`.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol.  
If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol.  
If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.  
If you select this option, set the following parameters:

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address:** indicates the IP address of the server where the backup data is stored.
- **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Server Shared Path:** indicates the backup path on the SFTP server.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **OBS:** indicates that backup files are stored in OBS.  
If you select this option, set the following parameters:
  - **Target Path:** indicates the OBS directory for storing backup data.
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.

 **NOTE**

Only MRS 3.1.0 or later supports data backup to OBS.

**Step 8** Click **OK**.

**Step 9** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version\_Data source\_Task execution time.tar.gz*.

----End

### 7.6.5.9 Backing Up HBase Service Data

#### Scenario

For HBase service data security, you need to back up HBase service data before a major operation on HBase (such as upgrade or migration). The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up HBase service data. Both automatic and manual backup tasks are supported.

The following situations may occur during the HBase service data backup:

- When a user creates an HBase table, **KEEP\_DELETED\_CELLS** is set to **false** by default. When the user backs up this HBase table, deleted data will be backed

up and junk data may exist after data restoration. This parameter can be set to **true** manually when an HBase table is created based on service requirements.

- When a user manually specifies the timestamp when writing data into an HBase table and the specified time is earlier than the last backup time of the HBase table, new data may not be backed up in incremental backup tasks.
- The HBase backup function cannot back up the access control lists (ACLs) for reading, writing, executing, creating, and managing HBase global or namespaces. After HBase data is restored, you need to reset the role permissions on FusionInsight Manager.
- If the backup data of the standby cluster is lost in an existing HBase backup task, the next incremental backup will fail, and you need to create an HBase backup task again. However, the next full backup task will be normal.

## Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
  - The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.
- Backup policies, including the backup task type, period, backup object, backup directory, and Yarn queue required by the backup task are planned based on service requirements.
- On the HDFS client, you have executed the **hdfs lsSnapshottableDir** command as user **hdfs** to check the list of directories for which HDFS snapshots have been created in the current cluster and ensured that the HDFS parent directory or subdirectory where data files to be backed up are stored does not have HDFS snapshots. Otherwise, the backup task cannot be created.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- The **fs.defaultFS** parameter settings of HBase are the same as those of Yarn and HDFS.

## Backing Up HBase Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.

**Manual** indicates that the backup task is executed manually.

**Table 7-27** Periodic backup parameters

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Started       | Indicates the time when the task is started for the first time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Period        | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Backup Policy | <ul style="list-style-type: none"> <li>• <b>Full backup at the first time and incremental backup subsequently</b></li> <li>• <b>Full backup every time</b></li> <li>• <b>Full backup once every n times</b></li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li> <li>• If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li> </ul> |

**Step 6** In **Configuration**, choose **HBase > HBase** under **Service data**.

**Step 7** Set **Path Type** of **HBase** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.



- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as `/hbase` or `/user/hbase/backup`.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory,

and the user group and owner group of the shared path must be **nobody:nobody**.)

- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.

If you select this option, set the following parameters:

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address:** indicates the IP address of the server where the backup data is stored.
- **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Server Shared Path:** indicates the backup path on the SFTP server.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.

**Step 8** Set **Maximum Number of Recovery Points** to the number of snapshots that can be retained in the cluster.

**Step 9** Set **Backup Content** to one or multiple HBase tables to be backed up.

You can select backup data using either of the following methods:

- Adding a backup data file

Click the name of a database in the navigation pane to show all the tables in the database, and select specified tables.

MRS 3.2.0 or later:

- a. Click **Add**.

- b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.
- c. Click **OK**.
- Selecting using regular expressions
  - a. Click **Query Regular Expression**.
  - b. Enter the namespace where the HBase tables are located in the first text box as prompted. The namespace must be the same as the existing namespace, for example, **default**.
  - c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all tables in the namespace, enter **([\s\S]\*?)**. To get tables whose names consist of letters and digits, for example, **tb 1**, enter **tb\d\***.
  - d. Click **Refresh** to view the displayed tables in **Directory Name**.
  - e. Click **Synchronize** to save the result.

 **NOTE**

- When entering regular expressions, click **+** or **-** to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.

**Step 10** Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The queue name is incorrect.
- The parent directory or subdirectory of the HDFS directory where HBase table data files to be backed up are stored has HDFS snapshots.
- The directory or table to be backed up does not exist.

**Step 11** Click **OK**.

**Step 12** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *xxx/Backup task name\_Data source\_Task creation time*, and the subdirectory is used to save latest data source backup files. All the backup file sets are stored in the related snapshot directories.

----End

## 7.6.5.10 Backing Up HDFS NameNode Data

### Scenario

To ensure NameNode service data security routinely or before a major operation on NameNode (such as upgrade or migration), you need to back up NameNode data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up NameNode data. Both automatic and manual backup tasks are supported.

## Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

## Backing Up HDFS NameNode Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.

**Manual** indicates that the backup task is executed manually.

**Table 7-28** Periodic backup parameters

| Parameter | Description                                                                               |
|-----------|-------------------------------------------------------------------------------------------|
| Started   | Indicates the time when the task is started for the first time.                           |
| Period    | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> . |

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Policy | <p>Only <b>Full backup every time</b> is supported.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li><li>If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li></ul> |

**Step 6** In **Configuration**, select **NameNode**.

**Step 7** Set **Path Type** of **NameNode** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files. By default, the backup files are stored in *Data storage path/LocalBackup/*.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select this option, set the following parameters:
  - **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Target NameNode IP Address**: indicates the service plane IP address of the NameNode in the standby cluster.
  - **Target Path**: indicates the path for storing backup files.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address:** indicates the IP address of the NAS server.
- **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the server where the backup data is stored.
  - **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username:** indicates the username for connecting to the server using the SFTP protocol.
  - **Password:** indicates the password for connecting to the server using the SFTP protocol.

- **Server Shared Path:** indicates the backup path on the SFTP server.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **OBS:** indicates that backup files are stored in OBS.  
If you select this option, set the following parameters:
  - **Target Path:** indicates the OBS directory for storing backup data.
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.

 **NOTE**

Only MRS 3.1.0 or later supports data backup to OBS.

**Step 8** Click **OK**.

**Step 9** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version\_Data source\_Task execution time.tar.gz*.

----End

## 7.6.5.11 Backing Up HDFS Service Data

### Scenario

To ensure HDFS service data security routinely or before a major operation on HDFS (such as upgrade or migration), you need to back up HDFS service data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up HDFS service data. Both automatic and manual backup tasks are supported.

 **NOTE**

Encrypted directories cannot be backed up or restored.

### Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager,

mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.

- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.
- Backup policies, including the backup task type, period, backup object, backup directory, and Yarn queue required by the backup task are planned based on service requirements.
- On the HDFS client, you have executed the `hdfs lsSnapshottableDir` command as user `hdfs` to check the list of directories for which HDFS snapshots have been created in the current cluster and ensured that the HDFS parent directory or subdirectory where data files to be backed up are stored does not have HDFS snapshots. Otherwise, the backup task cannot be created.
- If you want to back up data to NAS, you have deployed the NAS server in advance.

## Backing Up HDFS Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.

**Manual** indicates that the backup task is executed manually.

**Table 7-29** Periodic backup parameters

| Parameter | Description                                                                               |
|-----------|-------------------------------------------------------------------------------------------|
| Started   | Indicates the time when the task is started for the first time.                           |
| Period    | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> . |



| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Policy | <ul style="list-style-type: none"><li>● <b>Full backup at the first time and incremental backup subsequently</b></li><li>● <b>Full backup every time</b></li><li>● <b>Full backup once every n times</b></li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>● Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li><li>● If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li></ul> |

**Step 6** In **Configuration**, select **HDFS**.

**Step 7** Set **Path Type** of **HDFS** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

- **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.

- **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

If you select this option, set the following parameters:

- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address**: indicates the IP address of the server where the backup data is stored.
- **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username**: indicates the username for connecting to the server using the SFTP protocol.
- **Password**: indicates the password for connecting to the server using the SFTP protocol.
- **Server Shared Path**: indicates the backup path on the SFTP server.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

**Step 8** Set **Maximum Number of Recovery Points** to the number of snapshots that can be retained in the cluster.

**Step 9** Set **Backup Content** to one or multiple HDFS directories to be backed up based on service requirements.

You can select backup data using either of the following methods:

- Adding a backup data file  
Click the name of a database in the navigation pane to show all the tables in the database, and select specified tables.  
MRS 3.2.0 or later:
  - Click **Add**.
  - Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.
  - Click **OK**.
- Selecting using regular expressions

- a. Click **Query Regular Expression**.
- b. Enter the parent directory full path of the directory in the first text box as prompted. The directory must be the same as the existing directory, for example, **/tmp**.
- c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all files or subdirectories in the parent directory, enter **([\\s\\S]\*?)**. To get files whose names consist of letters and digits, for example, **file1**, enter **file\\d\***.
- d. Click **Refresh** to view the displayed directories in **Directory Name**.
- e. Click **Synchronize** to save the result.

 **NOTE**

- When entering regular expressions, click **+** or **-** to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.
- The backup directory cannot contain files that have been written for a long time. Otherwise, the backup task will fail. Therefore, you are not advised to perform operations on the top-level directory, such as **/user**, **/tmp**, and **/mr-history**.

**Step 10** Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The queue name is incorrect.
- The parent directory or subdirectory of the HDFS directory where data files to be backed up are stored has HDFS snapshots.
- The directory or table to be backed up does not exist.
- The name of the NameService is incorrect.

**Step 11** Click **OK**.

**Step 12** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Data source\_Task creation time*, and the subdirectory is used to save latest data source backup files. All the backup file sets are stored in the related snapshot directories.

----End

## 7.6.5.12 Backing Up Hive Service Data

### Scenario

To ensure Hive service data security routinely or before a major operation on Hive (such as upgrade or migration), you need to back up Hive service data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up Hive service data. Both automatic and manual backup tasks are supported.

- Hive backup and restoration cannot identify the service and structure relationships of objects such as Hive tables, indexes, and views. When executing backup and restoration tasks, you need to manage unified restoration points based on service scenarios to ensure proper service running.
- Hive backup and restoration do not support Hive on RDB data tables. You need to back up and restore original data tables in external databases independently.
- If the backup data of the standby cluster is lost in an existing Hive backup task that contains Hive on HBase tables, the next incremental backup will fail, and you need to create a Hive backup task again. However, the next full backup task will be normal.
- After the backup function of FusionInsight Manager is used to back up the HDFS directories at the Hive table level, the Hive tables cannot be deleted and recreated.

## Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
  - The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.
- Backup policies, including the backup task type, period, backup object, backup directory, and Yarn queue required by the backup task are planned based on service requirements.
- On the HDFS client, you have executed the `hdfs lsSnapshottableDir` command as user `hdfs` to check the list of directories for which HDFS snapshots have been created in the current cluster and ensured that the HDFS parent directory or subdirectory where data files to be backed up are stored does not have HDFS snapshots. Otherwise, the backup task cannot be created.
- If you want to back up data to NAS, you have deployed the NAS server in advance.

## Backing Up Hive Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.

**Manual** indicates that the backup task is executed manually.

**Table 7-30** Periodic backup parameters

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Started       | Indicates the time when the task is started for the first time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Period        | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Backup Policy | <ul style="list-style-type: none"><li>• <b>Full backup at the first time and incremental backup subsequently</b></li><li>• <b>Full backup every time</b></li><li>• <b>Full backup once every n times</b></li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li><li>• If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li></ul> |

**Step 6** In **Configuration**, choose **Hive > Hive**.

**Step 7** Set **Path Type** of **Hive** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select this option, set the following parameters:
  - **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.

- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as `/hbase` or `/user/hbase/backup`.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.

- **Password:** indicates the password set when the CIFS protocol is configured.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.

If you select this option, set the following parameters:

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address:** indicates the IP address of the server where the backup data is stored.
- **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Server Shared Path:** indicates the backup path on the SFTP server.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.

**Step 8** Set **Maximum Number of Recovery Points** to the number of snapshots that can be retained in the cluster.

**Step 9** Set **Backup Content** to one or multiple Hive tables to be backed up.



You can select backup data using either of the following methods:

- Adding a backup data file  
Click the name of a database in the navigation pane to show all the tables in the database, and select specified tables.  
MRS 3.2.0 or later:
  - a. Click **Add**.
  - b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.
  - c. Click **OK**.
- Selecting using regular expressions
  - a. Click **Query Regular Expression**.
  - b. Enter the database where the Hive tables are located in the first text box as prompted. The database must be the same as the existing database, for example, **default**.
  - c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all tables in the database, enter **([s\S]\*?)**. To get tables whose names consist of letters and digits, for example, **tb1**, enter **tb\d\***.
  - d. Click **Refresh** to view the displayed tables in **Directory Name**.
  - e. Click **Synchronize** to save the result.

 **NOTE**

- When entering regular expressions, click **+** or **-** to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.

**Step 10** Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The queue name is incorrect.
- The parent directory or subdirectory of the HDFS directory where data files to be backed up are stored has HDFS snapshots.
- The directory or table to be backed up does not exist.
- The name of the NameService is incorrect.

**Step 11** Click **OK**.

**Step 12** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Data source\_Task creation time*, and the subdirectory is used to save latest data source backup files. All the backup file sets are stored in the related snapshot directories.

----End

### 7.6.5.13 Backing Up IoTDB Metadata

#### Scenario

To ensure IoTDB metadata security and prevent the IoTDB service from being unavailable due to IoTDB metadata file damages, you need to back up IoTDB metadata. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up IoTDB metadata. Both automatic and manual backup tasks are supported.

#### Prerequisites

- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- Backup policies, including the backup task type, period, backup object, backup directory, and Yarn queue required by the backup task are planned based on service requirements.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.

#### Backing Up IoTDB Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.  
**Manual** indicates that the backup task is executed manually.

**Table 7-31** Periodic backup parameters

| Parameter | Description                                                                       |
|-----------|-----------------------------------------------------------------------------------|
| Started   | Indicates the time when the task is started for the first time.                   |
| Period    | The task execution interval. Value options include <b>Hours</b> and <b>Days</b> . |

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Policy | <p>Indicates a periodic backup policy.</p> <ul style="list-style-type: none"><li>● <b>Full backup at the first time and incremental backup subsequently</b></li><li>● <b>Full backup every time</b></li><li>● <b>Full backup once every n times</b></li></ul> <p><b>NOTE</b><br/>Incremental backup is not supported when component metadata is backed up. Only <b>Full backup every time</b> is supported.</p> |

**Step 6** In **Configuration**, select **IoTDB** under **Metadata and other data**.

**Step 7** Set **Path Type** of **IoTDB** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files. By default, the backup files are stored in *Data storage path/LocalBackup/*.  
If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.  
If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.  
If you select this option, set the following parameters:
  - **Destination NameService Name**: indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Destination Active NameNode IP Address**: indicates the service plane IP address of the active NameNode in the standby cluster.
  - **Destination Standby NameNode IP Address**: indicates the service plane IP address of the standby NameNode in the standby cluster.

- **Destination NameNode RPC Port:** indicates the value of `dfs.namenode.rpc.port` in the HDFS basic configuration of the standby cluster.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as `/hbase` or `/user/hbase/backup`.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol.  
If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.  
If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the server where the backup data is stored.
  - **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username:** indicates the username for connecting to the server using the SFTP protocol.
  - **Password:** indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path:** indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.

**Step 8** Click **OK**.

**Step 9** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version\_Data source\_Task execution time.tar.gz*.

----End

## 7.6.5.14 Backing Up IoTDB Service Data

### Scenario

To ensure IoTDB service data security routinely or before a major operation on IoTDB (such as upgrade or migration), you need to back up IoTDB service data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up IoTDB service data. Both automatic and manual backup tasks are supported.

### Prerequisites

To back up IoTDB service data to a remote HDFS, you need to meet the following conditions:

- A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
- The cluster where IoTDB is deployed must be in security mode.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.

## Backing Up IoTDB Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.  
**Manual** indicates that the backup task is executed manually.

**Table 7-32** Periodic backup parameters

| Parameter     | Description                                                                                                                                                                                                                                            |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Started       | Indicates the time when the task is started for the first time.                                                                                                                                                                                        |
| Period        | The task execution interval. Value options include <b>Hours</b> and <b>Days</b> .                                                                                                                                                                      |
| Backup Policy | Indicates a periodic backup policy. <ul style="list-style-type: none"><li>• <b>Full backup at the first time and incremental backup subsequently</b></li><li>• <b>Full backup every time</b></li><li>• <b>Full backup once every n times</b></li></ul> |

**Step 6** In **Configuration**, choose **IoTDB > IoTDB** under **Service data**.

**Step 7** Set **Path Type** of **IoTDB** to a backup directory type.

The following backup directory types are supported:

**RemoteHDFS:** indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name:** indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Destination Active NameNode IP Address:** indicates the service plane IP address of the active NameNode in the standby cluster.
- **Destination Standby NameNode IP Address:** indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Destination NameNode RPC Port:** indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the standby cluster.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.

**Step 8** Set **Backup Content** to one or multiple service data records to be backed up.

You can select backup data using either of the following methods:

- Adding a backup data file
  - a. Click **Add**.

- b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.
- c. Click **OK**.
- Selecting using regular expressions
  - a. Click **Query Regular Expression**.
  - b. Enter the parent directory full path of the directory in the first text box as prompted. The directory must be the same as the existing directory, for example, **/root**.
  - c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all files or subdirectories in the parent directory, enter **([\\s\\S]\*?)**. To get files whose names consist of letters and digits, for example, **file 1**, enter **file\\d\***.
  - d. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get objects containing **test**, enter **.\*test.\***. To get objects starting with **test**, enter **test.\***. To get objects ending with **test**, enter **.\*test**.
  - e. Click **Refresh** to view the displayed directories in **Directory Name**.
  - f. Click **Synchronize** to save the result.

**NOTE**

- When entering regular expressions, click **+** or **-** to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.
- The backup directory cannot contain files that have been written for a long time. Otherwise, the backup task will fail. Therefore, you are not advised to perform operations on the top-level directory, such as **/user**, **/tmp**, and **/mr-history**.

**Step 9** Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The data to be backed up does not exist.

**Step 10** Click **OK**.

**Step 11** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Data source\_Task creation time*, and the subdirectory is used to save latest data source backup files. All the backup file sets are stored in the related snapshot directories.

----End

## 7.6.5.15 Backing Up Kafka Metadata

### Scenario

To ensure Kafka metadata security or before a major operation on ZooKeeper (such as upgrade or migration), you need to back up Kafka metadata. The backup

data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up Kafka metadata. Both automatic and manual backup tasks are supported.

## Prerequisites

- To back up data to a remote HDFS, the following conditions must be met:
  - A standby cluster for backing up data has been created. The authentication mode must be the same as that of the active cluster.
  - If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
  - Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path/LocalBackup/* has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

## Backing Up Kafka Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** Click **Create**.

**Step 3** Set **Name** to the name of the backup task.

**Step 4** Select the cluster to be operated from **Backup Object**.

**Step 5** Set **Mode** to the type of the backup task.

**Periodic** indicates that the backup task is executed by the system periodically.

**Manual** indicates that the backup task is executed manually.

**Table 7-33** Periodic backup parameters

| Parameter | Description                                                                               |
|-----------|-------------------------------------------------------------------------------------------|
| Started   | Indicates the time when the task is started for the first time.                           |
| Period    | Indicates the task execution interval. The options include <b>Hours</b> and <b>Days</b> . |



| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Policy | <ul style="list-style-type: none"><li>● <b>Full backup at the first time and incremental backup subsequently</b></li><li>● <b>Full backup every time</b></li><li>● <b>Full backup once every n times</b></li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>● Incremental backup is not supported when Manager data and component metadata are backed up. Only <b>Full backup every time</b> is supported.</li><li>● If <b>Path Type</b> is set to <b>NFS</b> or <b>CIFS</b>, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.</li></ul> |

**Step 6** In **Configuration**, select **Kafka**.

**Step 7** Set **Path Type** of **Kafka** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files. By default, the backup files are stored in *Data storage path/LocalBackup/*.

If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.

- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

If you select this option, set the following parameters:

- **Target Path**: indicates the HDFS directory for storing the backup files. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address:** indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path:** indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path:** indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.
- **OBS:** indicates that backup files are stored in OBS. If you select this option, set the following parameters:

- **Target Path:** indicates the OBS directory for storing backup data.
- **Maximum Number of Backup Copies:** indicates the number of backup file sets that can be retained in the backup directory.

 **NOTE**

Only MRS 3.1.0 or later supports data backup to OBS.

**Step 8** Click **OK**.

**Step 9** In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name\_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version\_Data source\_Task execution time.tar.gz*.

----End

## 7.6.6 Restoring MRS Cluster Component Data

### 7.6.6.1 Restoring Manager Data (MRS2.x and Earlier)

#### Scenario

You need to restore metadata in the following scenarios: A user modifies or deletes data unexpectedly, data needs to be retrieved, system data becomes abnormal or does not achieve the expected result, all modules are faulty, and data is migrated to a new cluster.

This section describes how to restore metadata on MRS Manager. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the data that is generated after the data backup and before the data restoration will be lost.
  - Use the OMS data and LdapServer data backed up at the same time to restore data. Otherwise, the service and operation may fail.
  - By default, MRS clusters use DBService to store Hive metadata.
- 

#### Impact on the System

- After the data is restored, the data generated between the backup time and restoration time is lost.
- After the data is restored, the configuration of the components that depend on DBService may expire and these components need to be restarted.

## Prerequisites

- The data in the OMS and LdapServer backup files has been backed up at the same time.
- The status of the OMS resources and the LdapServer instances is normal. If the status is abnormal, data restoration cannot be performed.
- The status of the cluster hosts and services is normal. If the status is abnormal, data restoration cannot be performed.
- The cluster host topologies during data restoration and data backup are the same. If the topologies are different, data restoration cannot be performed and you need to back up data again.
- The services added to the cluster during data restoration and data backup are the same. If the services are different, data restoration cannot be performed and you need to back up data again.
- The status of the active and standby DBService instances is normal. If the status is abnormal, data restoration cannot be performed.
- The upper-layer applications depending on the MRS cluster have been stopped.
- On MRS Manager, you have stopped all the NameNode role instances whose data is to be recovered. Other HDFS role instances are running properly. After data is recovered, the NameNode role instances need to be restarted and cannot be accessed before the restart.
- You have checked whether NameNode backup files have been stored in the *Data save path/LocalBackup/* directory on the active management node.

## Restoring Manager Data

**Step 1** Check the location of backup data.

1. On MRS Manager, choose **System > Back Up Data**.
2. In the row where the specified backup task resides, choose **More > View History** in the **Operation** column to display the historical execution records of the backup task. In the window that is displayed, select a success record and click **View Backup Path** in the corresponding column to view its backup path information. Find the following information:
  - **Backup Object**: indicates the backup data source.
  - **Backup Path**: indicates the full path where backup files are stored.
3. Select the correct path, and manually copy the full path of backup files in **Backup Path**.

**Step 2** Create a restoration task.

1. On MRS Manager, choose **System > Recovery Management**.
2. On the page that is displayed, click **Create Restoration Task**.
3. Set **Task Name** to the name of the restoration task.

**Step 3** Select restoration sources.

In **Configuration**, select the metadata component whose data is to be restored.

**Step 4** Set the restoration parameters.

1. Set **Path Type** to a backup directory type.
2. The settings vary according to backup directory types:
  - **LocalDir**: indicates that the backup files are stored on the local disk of the active management node. If you select **LocalDir**, you need to set **Source Path** to specify the full path of the backup file. For example, *Data storage path/LocalBackup/Backup task name\_Task creation time/Data source\_Task execution time/Version number\_Data source\_Task execution time.tar.gz*.
  - **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster. If you select **SFTP**, set the following parameters:
    - **Source Path**: indicates the full HDFS path of a backup file. for example, *Backup path/Backup task name\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
    - **Source Instance Name**: indicates the name of NameService corresponding to the backup directory when a restoration task is being executed. The default value is **hacluster**.
3. Click **OK**.

**Step 5** Execute the restoration task.

In the restoration task list, locate the row where the created task resides, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and try to execute the task again by clicking **Start**.

**Step 6** Determine what metadata has been restored.

- If the OMS and LdapServer metadata is restored, go to [Step 7](#).
- If DBService data is restored, no further action is required.
- Restore NameNode data. On MRS Manager, choose **Services > HDFS > More > Restart Service**. The task is complete.

**Step 7** Restarting Manager for the recovered data to take effect

1. In MRS Manager, Choose **LdapServer > More > Restart Service** and click **OK**. Wait until the LdapServer service is restarted successfully.
2. Log in to the active management node. For details, see [Checking MRS Active/Standby Management Nodes](#).
3. Run the following command to restart OMS:

```
sh ${BIGDATA_HOME}/om-0.0.1/sbin/restart-oms.sh
```

The command has been executed successfully if the following information is displayed:

```
start HA successfully.
```

4. On MRS Manager, choose **KrbServer > More > Synchronize Configuration**. Do not select Restart the services and instances whose configuration has

expired. Click **OK** and wait until the KrbServer service configuration is synchronized and restarted successfully.

5. Choose **Services > More > Synchronize Configuration**. Do not select Restart the services and instances whose configuration has expired. Click **OK** and wait until the cluster is configured and synchronized successfully.
6. Choose **Services > More > Stop Cluster**. After the cluster is stopped, choose **Services > More > Start Cluster**.

----End

## 7.6.6.2 Restoring Manager Data (MRS 3.x and Later Versions)

### Scenario

Manager data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in FusionInsight Manager, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable.

System administrators can create a restoration task in FusionInsight Manager to recover Manager data. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that of data backup.
  - To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the Manager data that is generated after the data backup and before the data restoration will be lost.
- 

### Impact on the System

- In the restoration process, the Controller needs to be restarted and FusionInsight Manager cannot be logged in or operated during the restart.
- In the restoration process, all clusters need to be restarted and cannot be accessed during the restart.
- After data restoration, the data, such as system configuration, user information, alarm information, and audit information, that is generated after the data backup and before the data restoration will be lost. This may result in data query failure or cluster access failure.
- After the Manager data is recovered, the system forces the LdapServer of each cluster to synchronize data from the OLadp.

### Prerequisites

- To restore data from a remote HDFS, the following conditions must be met:
  - Prepare a standby cluster for restoring data, and ensure that data in this cluster has been backed up. For details, see [Backing Up Manager Data](#)

- (MRS 3.x and Later Versions)**. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, system mutual trust needs to be configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
  - Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
  - The status of the OMS resources and the LdapServer instances of each cluster is normal. If the status is abnormal, data restoration cannot be performed.
  - The status of the cluster hosts and services is normal. If the status is abnormal, data restoration cannot be performed.
  - The cluster host topologies during data restoration and data backup are the same. If the topologies are different, data restoration cannot be performed and you need to back up data again.
  - The services added to the cluster during data restoration and data backup are the same. If the topologies are different, data restoration cannot be performed and you need to back up data again.
  - The upper-layer applications that depend on the cluster are stopped.

## Restoring Manager Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of a specified task in the task list, choose **More > View History** to view the historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.  
Select the correct item, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restore > Restoring Management**. On the displayed page, click **Create**.

**Figure 7-23** Creating a restoration task

Restoration Management > Create Restoration Task

• Task Name:  The task name contains 3 to 128 characters, including digits, letters, and underscores (\_), and cannot be empty.

• Recovery Object:

• Restoration Configuration: Metadata and other data

- DBService
- NameNode (The NameNode instances must be stopped before the restoration.)
- Yarn
- HBase
- Kafka

Service data

- HDFS
- HBase
- Hive

**Step 4** Set **Task Name** to the name of the restoration task.

**Step 5** Set **Recovery Object** to **OMS**.

**Step 6** Select **OMS**.

**Step 7** Set **Path Type** of **OMS** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.  
If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Version\_Data source\_Task execution time.tar.gz*.
- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.  
If you select **LocalHDFS**, set the following parameters:
  - **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/Backup task name\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **Cluster for Restoration**: Enter the name of the cluster used during restoration task execution.
  - **Source NameService Name**: indicates the NameService name that corresponds to the backup directory when a restoration task is executed. The default value is **hacluster**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.  
If you select **RemoteHDFS**, set the following parameters:
  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the



- remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **Source Cluster**: Select the cluster of the Yarn queue used by the recovery data.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol. If you select **NFS**, set the following parameters:
    - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
    - **Server IP Address**: indicates the IP address of the NAS server.
    - **Source Path**: indicates the complete path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
    - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
    - **Server IP Address**: indicates the IP address of the NAS server.
    - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
    - **Username**: indicates the username set when the CIFS protocol is configured.
    - **Password**: indicates the password set when the CIFS protocol is configured.
    - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **SFTP**: indicates that backup files are stored in the server using the SFTP protocol. If you select **SFTP**, set the following parameters:
    - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.

- **Server IP Address:** indicates the IP address of the server where the backup data is stored.
- **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Source Path:** indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **OBS:** indicates that backup files are stored in OBS.  
If you select **OBS**, set the following parameters:
  - **Source Path:** indicates the full OBS path of a backup file, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.

 **NOTE**

Only MRS 3.1.0 or later supports saving backup files in OBS.

**Step 8** Click **OK**.

**Step 9** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

**Step 10** Log in to the active and standby management nodes as user **omm**.

**Step 11** Run the following command to restart OMS:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/restart-oms.sh
```

The command is run successfully if the following information is displayed:

```
start HA successfully.
```

Run `sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh` to check whether **HAAllResOK** of the management node is **Normal** and whether FusionInsight Manager can be logged in again. If yes, OMS is restarted successfully.

**Step 12** On FusionInsight Manager, click **Cluster**, click the name of the target cluster, and choose **Services > KrbServer**. On the displayed page, choose **More > Synchronize Configuration**, click **OK**, and wait for the KrbServer configuration to be synchronized and the service to be restarted.

**Step 13** Choose **Cluster**, click the name of the desired cluster, and choose **More > Synchronize Configurations**, click **OK**, and wait until the cluster configuration is synchronized successfully.

- Step 14** On FusionInsight Manager, click **Cluster**, click the name of the target cluster, and choose **More > Restart**. On the displayed page, enter the password of the current login user, click **OK**, and wait for the cluster to be restarted.

----End

### 7.6.6.3 Restoring CDL Service Data

#### Scenario

CDL data needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After an administrator performs major operations (such as upgrade and data adjustment) on CDL, an exception occurs or the expected result is not achieved. All modules are faulty and become unavailable. Data is migrated to a new cluster.

CDL metadata is stored in DBService and Kafka. A system administrator can create DBService and Kafka restoration tasks on FusionInsight Manager to restore CDL data. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the DBService and Kafka data that is generated after the data backup and before the data restoration will be lost.
  - By default, MRS clusters use DBService to store metadata of Hive, Hue, Loader, Spark, CDL, and Oozie. Restoring DBService data will restore the metadata of all these components.
- 

#### Impact on the System

- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is restored, the configurations of the components that depend on DBService may expire and these components need to be restarted.
- After the metadata is restored, the offset information stored on ZooKeeper by Kafka consumers is rolled back, resulting in repeated consumption.

#### Prerequisites

- To restore data from a remote HDFS, the following conditions must be met:
  - A standby cluster for restoring data has been created, and data in this cluster has been backed up. For details, see [Backing Up CDL Service Data](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.

- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The status of the active and standby DBService instances is normal. If the status is abnormal, data restoration cannot be performed.
- The Kafka service has been stopped. After the restoration is complete, start the Kafka service.

## Restoring CDL Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the row where the specified backup task is located, choose **More > View History** in the **Operation** column to display the historical execution records of the backup task.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object:** indicates the backup data source.
- **Backup Path:** indicates the full path where backup files are stored.  
Select the correct path, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In the **Restoration Configuration** area, select **DBService** and **Kafka**.

**Step 8** Set **Path Type** of **DBService** to a backup directory type. For details about how to configure the parameters, see [Step 8](#).

**Step 9** Set **Path Type** of **Kafka** to a backup directory type. For details about how to configure the parameters, see [Step 8](#).

**Step 10** Click **OK**.

**Step 11** In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

## 7.6.6.4 Restoring ClickHouse Metadata

### Scenario

ClickHouse metadata needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After a user performs major operations (such as upgrade and migration) on ClickHouse, an exception occurs or the expected result is not achieved. The ClickHouse component is faulty and becomes unavailable. Data is migrated to a new cluster.

Users can create a ClickHouse restoration task on FusionInsight Manager. Only manual restoration tasks are supported.

---

#### NOTICE

- This function is supported only by MRS 3.1.0 or later.
  - Data restoration can be performed only when the system version is consistent with that during data backup.
  - To restore ClickHouse metadata when the service is running properly, you are advised to manually back up the latest ClickHouse metadata before restoration. Otherwise, the ClickHouse metadata that is generated after the data backup and before the data restoration will be lost.
  - ClickHouse metadata restoration and service data restoration cannot be performed at the same time. Otherwise, service data restoration fails. You are advised to restore service data after metadata restoration is complete.
- 

### Impact on the System

- After the metadata is restored, the data generated after the data backup and before the data restoration is lost.
- After the metadata is restored, the ClickHouse upper-layer applications need to be started.

### Prerequisites

- You have checked the path for storing ClickHouse metadata backup files.
- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up ClickHouse Metadata](#). If the active and standby clusters are deployed in security mode and they are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active and standby clusters are deployed in normal mode, no mutual trust is required.
- In an active/standby cluster, the value of `HADOOP_RPC_PROTECTION` of ClickHouse must be the same as that of `hadoop.rpc.protection` in the HDFS when you restore data from the remote HDFS to the local host.

## Restoring ClickHouse Metadata

- Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of the specified task in the task list, choose **More > View History**.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object:** indicates the backup data source.
- **Backup Path:** indicates the full path where backup files are stored.  
Select the correct path, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In **Restoration Configuration**, select **ClickHouse** under **Metadata and other data**.

**Step 8** Set **Path Type** of **ClickHouse** to a restoration directory type.

The configurations vary based on backup directory types:

- **LocalDir:** indicates that data is restored from the local disk of the active management node.  
If you select this value, you also need to configure the following parameters:
  - **Source Path:** backup file to be restored, for example, *Backup task name\_Data source\_Task execution time.tar.gz*.
- **RemoteHDFS:** indicates that data is restored from the HDFS directory of the standby cluster.

If you select this option for MRS 3.2.0 or later clusters, you also need to configure the following parameters:

- **Source NameService Name:** indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source Active NameNode IP Address:** indicates the service plane IP address of the active NameNode in the standby cluster.
- **Source Standby NameNode IP Address:** indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Source NameNode RPC Port:** indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the destination cluster.
- **Source Path:** indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Data source\_Task execution time.tar.gz*.

- **Logical Cluster** of MRS 3.2.0 or later: Enter the ClickHouse logical cluster whose data has been backed up.

If you select this option for MRS 3.1.0 or 3.1.2 clusters, you also need to configure the following parameters:

- **Source NameService Name:** indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address:** indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Source Path:** indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Data source\_Task execution time.tar.gz*.

**Step 9** Click **OK**.

**Step 10** In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column. In the displayed dialog box, click **OK** to start the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

**Step 11** Choose **Cluster > Services** and start the ClickHouse service.

----End

## 7.6.6.5 Restoring ClickHouse Service Data

### Scenario

ClickHouse data needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After a user performs major operations (such as upgrade and migration) on ClickHouse, an exception occurs or the expected result is not achieved. All modules are faulty and become unavailable. Data is migrated to a new cluster.

Users can create a ClickHouse restoration task on FusionInsight Manager to restore data. Only manual restoration tasks are supported.

The ClickHouse backup and restoration functions cannot identify the service and structure relationships of objects such as ClickHouse tables, indexes, and views. When executing backup and restoration tasks, you need to manage unified restoration points based on service scenarios to ensure proper service running.

**NOTICE**

- This function is supported only by MRS 3.1.0 or later.
- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the ClickHouse data that is generated after the data backup and before the data restoration will be lost.
- ClickHouse metadata restoration and service data restoration cannot be performed at the same time. Otherwise, service data restoration fails. You are advised to restore service data after metadata restoration is complete.

## Impact on the System

- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is restored, the ClickHouse upper-layer applications need to be started.

## Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up ClickHouse Service Data](#). If the active and standby clusters are deployed in security mode and they are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active and standby clusters are deployed in normal mode, no mutual trust is required.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The database for storing restored data tables, the HDFS save path of data tables, and the list of users who can access restored data are planned.
- The ClickHouse backup file save path is correct.
- The ClickHouse upper-layer applications are stopped.
- In an active/standby cluster, the value of `HADOOP_RPC_PROTECTION` of ClickHouse must be the same as that of `hadoop.rpc.protection` in the HDFS when you restore data from the remote HDFS to the local host.

## Restoring ClickHouse Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the row where the specified backup task is located, choose **More > View History** in the **Operation** column to display the historical execution records of the backup task.



In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object:** indicates the backup data source.
- **Backup Path:** indicates the full path where backup files are stored.  
Select the correct path, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In **Restoration Configuration**, select **ClickHouse** under **Service data**.

**Step 8** Set **Path Type** of **ClickHouse** to a backup directory type.

Currently, the backup directory supports only the **RemoteHDFS** type.

- **RemoteHDFS:** indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option for MRS 3.2.0 or later clusters, you also need to configure the following parameters:

- **Source NameService Name:** indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source Active NameNode IP Address:** indicates the service plane IP address of the active NameNode in the standby cluster.
- **Source Standby NameNode IP Address:** indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Source NameNode RPC Port:** indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the destination cluster.
- **Source Path:** indicates the full path of the HDFS directory for storing backup data of the standby cluster. For details, see **Backup Path** obtained in **Step 2**. for example, *Backup path/Backup task name\_Data source\_Task creation time*.

If you select this option for MRS 3.1.0 or 3.1.2 clusters, you also need to configure the following parameters:

- **Source NameService Name:** indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.

- **Source NameNode IP Address:** indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Source Path:** indicates the full path of HDFS directory for storing backup data of the standby cluster. For details, see the **Backup Path** obtained in step [Step 2](#), for example, *Backup path/Backup task name\_Data source\_Task creation time*.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.

**Step 9** Click **OK**.

**Step 10** In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

### 7.6.6.6 Restoring DBService Metadata

#### Scenario

DBService data needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in DBService, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover DBService data. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the DBService data that is generated after the data backup and before the data recovery will be lost.
  - By default, MRS clusters use DBService to store metadata of Hive, Hue, Loader, Spark, CDL, and Oozie. Restoring DBService data will restore the metadata of all these components.
-

## Impact on the System

- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is restored, the configurations of the components that depend on DBService may expire and these components need to be restarted.

## Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up DBService Data](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The status of the active and standby DBService instances is normal. If the status is abnormal, data restoration cannot be performed.

## Restoring DBService Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of a specified task in the task list, choose **More > View History** to view the historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.  
Select the correct item, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In the **Restoration Configuration** area, select **DBService**.

**Step 8** Set **Path Type** of **DBService** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.  
If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Version\_Data source\_Task execution time.tar.gz*.
- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.  
If you select **LocalHDFS**, set the following parameters:
  - **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/Backup task name\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **Source NameService Name**: indicates the NameService name that corresponds to the backup directory when a restoration task is executed. The default value is **hacluster**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.  
If you select **RemoteHDFS**, set the following parameters:
  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.  
If you select **NFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol.  
If you select **CIFS**, set the following parameters:

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address:** indicates the IP address of the NAS server.
- **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
- **Username:** indicates the username set when the CIFS protocol is configured.
- **Password:** indicates the password set when the CIFS protocol is configured.
- **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.

- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.

If you select **SFTP**, set the following parameters:

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address:** indicates the IP address of the server where the backup data is stored.
- **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Source Path:** indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.

- **OBS:** indicates that backup files are stored in OBS.

If you select **OBS**, set the following parameters:

- **Source Path:** indicates the full OBS path of a backup file, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.

#### NOTE

Only MRS 3.1.0 or later supports saving backup files in OBS.

**Step 9** Click **OK**.

**Step 10** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.

- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

### 7.6.6.7 Restoring Flink Metadata

#### Scenario

Flink metadata needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After an administrator performs major operations (such as upgrade and data adjustment) on Flink, an exception occurs or the expected result is not achieved. The Flink component is faulty and becomes unavailable. Data is migrated to a new cluster.

System administrators can create a Flink restoration task on FusionInsight Manager. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To restore Flink metadata when the service is running properly, you are advised to manually back up the latest Flink metadata before restoration. Otherwise, the Flink metadata that is generated after the data backup and before the data restoration will be lost.
  - Flink metadata restoration and service data restoration cannot be performed at the same time. Otherwise, service data restoration fails. You are advised to restore service data after metadata restoration is complete.
- 

#### Impact on the System

- Before restoring the metadata, you need to stop the Flink service. During this period, all upper-layer applications are affected and cannot work properly.
- After the metadata is restored, the data generated after the data backup and before the data restoration is lost.
- After the metadata is restored, the Flink upper-layer applications of Solr need to be started.

#### Prerequisites

- You have checked the path for storing Flink metadata backup files.
- The Flink service has been stopped before its metadata is restored.
- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up Flink Metadata](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.

- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).

## Restoring Flink Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of the specified task in the task list, choose **More > View History**.

In the displayed window, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved. Select the correct path, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In **Restoration Configuration**, select **Flink** under **Metadata and other data**.

**Step 8** Set **Path Type** of **Flink** to a restoration directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that data is restored from the local disk of the active management node.  
If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Backup task name\_Data source\_Task execution time.tar.gz*.
- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.  
If you select **LocalHDFS**, set the following parameters:
  - **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/Backup task name\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **Source NameService Name**: indicates the NameService name that corresponds to the backup directory when a restoration task is executed.
- **RemoteHDFS**: indicates that data is restored from the HDFS directory of the standby cluster.  
If you select **RemoteHDFS**, set the following parameters:
  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**,

**haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.

- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Data source\_Task execution time.tar.gz*.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.

**Step 9** Click **OK**.

**Step 10** In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column. In the displayed dialog box, click **OK** to start the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

**Step 11** Choose **Cluster > Services** and start the Flink service.

----End

## 7.6.6.8 Restoring HBase Metadata

### Scenario

To ensure HBase metadata security (including tableinfo files and HFiles) or before a major operation on HBase system tables (such as upgrade or migration), you need to back up HBase metadata to prevent HBase service unavailability caused by HBase system table directory or file damages. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

System administrators can create a recovery task in FusionInsight Manager to recover HBase metadata. Only manual restoration tasks are supported.



**NOTICE**

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the HBase data that is generated after the data backup and before the data recovery will be lost.
- It is recommended that a data restoration task restore the metadata of only one component to prevent the data restoration of other components from being affected by stopping a service or instance. If data of multiple components is restored at the same time, data restoration may fail.  
HBase metadata cannot be restored at the same time as NameNode metadata. As a result, data restoration fails.

## Impact on the System

- Before restoring the metadata, you need to stop the HBase service, during which the HBase upper-layer applications are unavailable.
- After the metadata is restored, the data generated after the data backup and before the data restoration is lost.
- After the metadata is restored, the upper-layer applications of HBase need to be started.

## Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up HBase Metadata](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- You have checked the path for storing HBase metadata backup files.
- The HBase service has been stopped before its metadata is restored.

## Restoring HBase Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of a specified task in the task list, choose **More > View History** to view the historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.

- **Backup Path** specifies the full path where the backup files are saved.  
Select the correct item, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In **Restoration Configuration**, select **HBase** under **Metadata and other data**.

**Step 8** Set **Path Type** of **HBase** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.  
If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Version\_Data source\_Task execution time.tar.gz*.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select **RemoteHDFS**, set the following parameters:

- **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.  
If you select **NFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.

- **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **CIFS:** indicates that backup files are stored in NAS using the CIFS protocol.  
If you select **CIFS**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.  
If you select **SFTP**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the server where the backup data is stored.
  - **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username:** indicates the username for connecting to the server using the SFTP protocol.
  - **Password:** indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path:** indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **OBS:** indicates that backup files are stored in OBS.  
If you select **OBS**, set the following parameters:
  - **Source Path:** indicates the full OBS path of a backup file, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.

 **NOTE**

Only MRS 3.1.0 or later supports saving backup files in OBS.

**Step 9** Click **OK**.

**Step 10** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

### 7.6.6.9 Restoring HBase Service Data

#### Scenario

HBase data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in HBase, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover HBase data. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the HBase data that is generated after the data backup and before the data recovery will be lost.
- 

#### Impact on the System

- During the data recovery process, the system disables the HBase table to be recovered and the table cannot be accessed in this moment. The data recovery process takes several minutes, during which the HBase upper-layer applications are unavailable.
- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is recovered, the HBase upper-layer applications need to be started.

#### Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up HBase Service Data](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust](#)

**Between MRS Clusters.** If the active cluster is deployed in normal mode, no mutual trust is required.

- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The directory for saving the backup file has been checked.
- The HBase upper-layer applications have been stopped.

## Restoring HBase Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of a specified task in the task list, choose **More > View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.  
Select the correct item, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In **Restoration Configuration**, select **HBase** under **Service Data**.

**Step 8** Set **Path Type** of **HBase** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select **RemoteHDFS**, set the following parameters:
  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.

- **Source Path:** indicates the full path of the backup file in the HDFS, for example, *Backup path/xxx/Backup task name\_Data source\_Task creation time*. To obtain the backup path, locate the row that contains the target backup task in the backup management list, click **More > View History** in the **Operation** column, and click **View** in the **Backup Path** column.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution.
- **Recovery Point List:** Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol. If you select **NFS**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/xxx/Backup task name\_Data source\_Task creation time*. To obtain the backup path, locate the row that contains the target backup task in the backup management list, click **More > View History** in the **Operation** column, and click **View** in the **Backup Path** column.
  - **Queue Name:** indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List:** Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/xxx/Backup task name\_Data source\_Task*

*creation time*. To obtain the backup path, locate the row that contains the target backup task in the backup management list, click **More > View History** in the **Operation** column, and click **View** in the **Backup Path** column.

- **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
- **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
- **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

If you select **SFTP**, set the following parameters:

- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address**: indicates the IP address of the server where the backup data is stored.
- **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username**: indicates the username for connecting to the server using the SFTP protocol.
- **Password**: indicates the password for connecting to the server using the SFTP protocol.
- **Source Path**: indicates the full path of the backup file on the backup server, for example, *Backup path/xxx/Backup task name\_Data source\_Task creation time*. To obtain the backup path, locate the row that contains the target backup task in the backup management list, click **More > View History** in the **Operation** column, and click **View** in the **Backup Path** column.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
- **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
- **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

**Step 9** Set **Backup Data** column in **Data Configuration** to one or multiple backup data sources to be recovered. In the **Target Namespace** column, specify the target naming space after backup data recovery.

You are advised to set **Target Namespace** to a location that is different from the backup naming space.

**Step 10** Set **Force recovery** to **true**, which indicates to forcibly recover all backup data when a data table with the same name already exists. If the data table contains

new data added after backup, the new data will be lost after the data recovery. If you set the parameter to **false**, the restoration task is not executed if a data table with the same name exists.

**Step 11** Click **Verify** to check whether the restoration task is configured correctly.

- If the queue name is incorrect, the verification fails.
- If the specified naming space does not exist, the verification fails.
- If the forcible overwrite conditions are not met, the verification fails.

**Step 12** Click **OK** to save the settings.

**Step 13** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

**Step 14** Check whether HBase data is restored in an environment where HBase is newly installed or reinstalled.

- If yes, the administrator needs to set new permission for roles on FusionInsight Manager based on the original service plan.
- If no, no further operation is required.

----End

### 7.6.6.10 Restoring HDFS NameNode Metadata

#### Scenario

NameNode data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in NameNode, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover NameNode data. Only manual restoration tasks are supported.



**NOTICE**

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the NameNode data that is generated after the data backup and before the data recovery will be lost.
- It is recommended that a data restoration task restore the metadata of only one component to prevent the data restoration of other components from being affected by stopping a service or instance. If data of multiple components is restored at the same time, data restoration may fail.  
HBase metadata cannot be restored at the same time as NameNode metadata. As a result, data restoration fails.

## Impact on the System

- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is recovered, the NameNode needs to be restarted and is unavailable during the restart.
- After data is restored, metadata and service data may not be matched, the HDFS enters the security mode, and the HDFS service fails to be started. .

## Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up HDFS NameNode Data](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- On FusionInsight Manager, all the NameNode role instances whose data is to be recovered are stopped. Other HDFS role instances must keep running. After data is recovered, the NameNode role instances need to be restarted. The NameNode role instances cannot be accessed during the restart.
- The NameNode backup files are stored *Data path/LocalBackup/* on the active management node.

## Restoring HDFS NameNode Metadata

- Step 1** On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and choose **Services > HDFS**. On the displayed page, click **Instances** and click **NameNode** to check whether the NameNode instances of the data to be restored

are stopped. If the NameNode instances are not stopped, click **Stop Instance** in the upper right corner.

**Step 2** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 3** In the **Operation** column of a specified task in the task list, choose **More > View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.  
Select the correct item, and manually copy the full path of backup files in **Backup Path**.

**Step 4** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 5** Click **Create**.

**Step 6** Set **Task Name** to the name of the restoration task.

**Step 7** Select the cluster to be operated from **Recovery Object**.

**Step 8** In the **Restoration Configuration** area, select **NameNode**.

**Step 9** Set **Path Type** of **NameNode** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.  
If you select **LocalDir**, set the following parameters:
  - **Source Path**: indicates the full path of the backup file on the local disk, for example, *Backup path/Backup task name\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select **RemoteHDFS**, set the following parameters:

- **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.

- **Source Path:** indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **NFS:** indicates that backup files are stored in the NAS using the NFS protocol. If you select **NFS**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **CIFS:** indicates that backup files are stored in the NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
  - **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol. If you select **SFTP**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the server where the backup data is stored.
  - **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.

- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Source Path:** indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **OBS:** indicates that backup files are stored in OBS.

If you select **OBS**, set the following parameters:

- **Source Path:** indicates the full OBS path of a backup file, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.

#### NOTE

Only MRS 3.1.0 or later supports saving backup files in OBS.

**Step 10** Click **OK**.

**Step 11** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

**Step 12** On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and choose **Services > HDFS**. On the displayed page, click **Configurations** and click **All Configurations**.

On the displayed page, enter the password of the administrator who has logged in for authentication and click **OK**. After the system displays "Operation succeeded", click **Finish**. The service is started successfully.

----End

## 7.6.6.11 Restoring HDFS Service Data

### Scenario

HDFS data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in the HDFS, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover HDFS data. Only manual restoration tasks are supported.

**NOTICE**

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the HDFS data that is generated after the data backup and before the data recovery will be lost.
- The HDFS restoration operation cannot be performed for the directories used by running Yarn tasks, for example, `/tmp/logs`, `/tmp/archived`, and `/tmp/hadoop-yarn/staging`. Otherwise, data restoration using Distcp tasks fails due to file loss.

## Impact on the System

- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is recovered, the HDFS upper-layer applications need to be started.

## Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up HDFS Service Data](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The HDFS backup file save path is correct.
- The HDFS upper-layer applications are stopped.

## Restoring HDFS Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of a specified task in the task list, choose **More > View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.

- **Backup Path** specifies the full path where the backup files are saved.  
Select the correct item, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In **Restoration Configuration**, select **HDFS** under **Service Data**.

**Step 8** Set **Path Type** of **HDFS** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select **RemoteHDFS**, set the following parameters:

- **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **NFS**: indicates that backup files are stored in NAS using the NFS protocol. If you select **NFS**, set the following parameters:
    - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.

- **Server IP Address:** indicates the IP address of the NAS server.
- **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time*.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution.
- **Recovery Point List:** Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
- **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **CIFS:** indicates that backup files are stored in NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time*.
  - **Queue Name:** indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List:** Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.  
If you select **SFTP**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the server where the backup data is stored.

- **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Source Path:** indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name\_Data source\_Task creation time*.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution.
- **Recovery Point List:** Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
- **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.

**Step 9** In the **Backup Data** column of the **Data Configuration** page, select one or more pieces of backup data that needs to be restored based on service requirements. In the **Target Path** column, specify the target location after backup data restoration.

You are advised to set **Target Path** to a new path that is different from the backup path.

**Step 10** Click **Verify** to check whether the restoration task is configured correctly.

- If the queue name is incorrect, the verification fails.
- If the specified directory to be restored does not exist, the verification fails.

**Step 11** Click **OK**.

**Step 12** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

## 7.6.6.12 Restoring Hive Service Data

### Scenario

Hive data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in the Hive, an exception occurs or the operation has not



achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover Hive data. Only manual restoration tasks are supported.

Hive backup and restoration cannot identify the service and structure relationships of objects such as Hive tables, indexes, and views. When executing backup and restoration tasks, you need to manage unified restoration points based on service scenarios to ensure proper service running.

---

**NOTICE**

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the Hive data that is generated after the data backup and before the data recovery will be lost.
- 

## Impact on the System

- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is recovered, the Hive upper-layer applications need to be started.

## Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up Hive Service Data](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The database for storing restored data tables, the HDFS save path of data tables, and the list of users who can access restored data are planned.
- The Hive backup file save path is correct.
- The Hive upper-layer applications are stopped.

## Restoring Hive Service Data

- Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of a specified task in the task list, choose **More > View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.  
Select the correct item, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In the **Restoration Configuration** area, select **Hive**.

**Step 8** Set **Path Type** of **Hive** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select **RemoteHDFS**, set the following parameters:
  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select a Hive backup file set that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

- **NFS:** indicates that backup files are stored in NAS using the NFS protocol. If you select **NFS**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time*.
  - **Queue Name:** indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List:** Click **Refresh** and select a Hive backup file set that has been backed up in the standby cluster.
  - **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **CIFS:** indicates that backup files are stored in NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time*.
  - **Queue Name:** indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List:** Click **Refresh** and select a Hive backup file set that has been backed up in the standby cluster.
  - **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **100**.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.  
If you select **SFTP**, set the following parameters:

- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address:** indicates the IP address of the server where the backup data is stored.
- **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username:** indicates the username for connecting to the server using the SFTP protocol.
- **Password:** indicates the password for connecting to the server using the SFTP protocol.
- **Source Path:** indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name\_Data source\_Task creation time*.
- **Queue Name:** indicates the name of the Yarn queue used for backup task execution.
- **Recovery Point List:** Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
- **Target NameService Name:** indicates the NameService name of the backup directory. The default value is **hacluster**.
- **Maximum Number of Maps:** indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s):** indicates the maximum bandwidth of a map. The default value is **1**.

**Step 9** Set **Backup Data** in the **Data Configuration** to one or multiple backup data sources to be recovered based on service requirements. In the **Target Database** and **Target Path** columns, specify the target database and file save path after backup data recovery.

Configuration restrictions:

- Data can be restored to the original database, but data tables must be stored in a new path that is different from the backup path.
- To restore Hive index tables, select the Hive data tables that correspond to the Hive index tables to be restored.
- If a new restoration directory is selected to avoid affecting the current data, HDFS permission must be manually granted so that users who have permission of backup tables can access this directory.
- Data can be restored to other databases. In this case, HDFS permission must be manually granted so that users who have permission of backup tables can access the HDFS directory that corresponds to the database.

**Step 10** Set **Force recovery** to **true**, which indicates to forcibly recover all backup data when a data table with the same name already exists. If the data table contains new data added after backup, the new data will be lost after the data recovery. If you set the parameter to **false**, the restoration task is not executed if a data table with the same name exists.

**Step 11** Click **Verify** to check whether the restoration task is configured correctly.

- If the queue name is incorrect, the verification fails.
- If the specified directory to be restored does not exist, the verification fails.
- If the forcible overwrite conditions are not met, the verification fails.

**Step 12** Click **OK**.

**Step 13** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

### 7.6.6.13 Restoring IoTDB Metadata

#### Scenario

To ensure IoTDB metadata security and prevent the IoTDB service from being unavailable due to IoTDB file damage, IoTDB metadata needs to be backed up. In this way, the system can restore data timely when an exception is reported or an operation does not achieve the expected result, minimizing the impact on services.

System administrators can create an IoTDB restoration task on FusionInsight Manager. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the IoTDB data that is generated after the data backup and before the data restoration will be lost.
  - You are advised to restore the metadata of only one component in a restoration task to prevent the stop of a service or instance from affecting the data restoration of other components. If data of multiple components is restored at the same time, data restoration may fail.
- 

#### Impact on the System

After the metadata is restored, the data generated after the data backup and before the data restoration is lost.

#### Restoring IoTDB Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the row where the specified backup task is located, choose **More > View History** in the **Operation** column to display the historical execution records of the backup task.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object:** indicates the backup data source.
- **Backup Path:** indicates the full path where backup files are stored.  
Select the correct path, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In **Restoration Configuration**, select **IoTDB** under **Metadata and other data**.

**Step 8** Select a backup directory type for **Path Type**.

The configurations vary based on backup directory types:

- **LocalDir:** indicates that the backup files are stored on the local disk of the active management node.  
If you select this option, you also need to set **Source Path**, which indicates the backup file to be restored, for example, *Version\_Data source\_Task execution time.tar.gz*.
- **NFS:** indicates that backup files are stored in NAS using the NFS protocol.  
If you select this option, configure the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Source Path:** indicates the complete path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **RemoteHDFS:** indicates that the backup files are stored in the HDFS directory of the standby cluster.  
If you select this option, configure the following parameters:
  - **Source NameService Name:** indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source Active NameNode IP Address:** indicates the service plane IP address of the active NameNode in the standby cluster.

- **Source Standby NameNode IP Address:** indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Source NameNode RPC Port:** indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the standby cluster.
- **Source Path:** indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **CIFS:** indicates that backup files are stored in NAS using the CIFS protocol.  
If you select this option, configure the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the NAS server.
  - **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username:** indicates the username set when the CIFS protocol is configured.
  - **Password:** indicates the password set when the CIFS protocol is configured.
  - **Source Path:** indicates the complete path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **SFTP:** indicates that backup files are stored in the server using the SFTP protocol.  
If you select this option, configure the following parameters:
  - **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address:** indicates the IP address of the server where the backup data is stored.
  - **Port:** indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username:** indicates the username for connecting to the server using the SFTP protocol.
  - **Password:** indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path:** indicates the complete path of the backup file on the backup server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.

**Step 9** Click **OK**.

**Step 10** In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.

- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

**Step 11** Choose **Cluster > Services** and start the IoTDB service.

----End

## 7.6.6.14 Restoring IoTDB Service Data

### Scenario

IoTDB service data needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After an administrator performs major operations (such as upgrade and data adjustment) on IoTDB, an exception occurs or the expected result is not achieved. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create an IoTDB restoration task on FusionInsight Manager. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the IoTDB data that is generated after the data backup and before the data restoration will be lost.
- 

### Impact on the System

- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is restored, the IoTDB upper-layer applications need to be started.

### Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up IoTDB Service Data](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.



- The IoTDB backup file save path is correct.
- The IoTDB upper-layer applications are stopped.

## Restoring IoTDB Service Data

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the row where the specified backup task is located, choose **More > View History** in the **Operation** column to display the historical execution records of the backup task.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object:** indicates the backup data source.
- **Backup Path:** indicates the full path where backup files are stored.  
Select the correct path, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In **Restoration Configuration**, choose **IoTDB > IoTDB** under **Service Data**.

**Step 8** Set **Path Type** of **IoTDB** to a backup directory type.

The following backup directory types are supported:

**RemoteHDFS:** indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, configure the following parameters:

- **Source NameService Name:** indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode:** indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source Active NameNode IP Address:** indicates the service plane IP address of the active NameNode in the standby cluster.
- **Source Standby NameNode IP Address:** indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Source NameNode RPC Port:** indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the standby cluster.
- **Source Path:** indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time*.

- **Recovery Point List:** Click **Refresh** and select an IoTDB directory that has been backed up in the standby cluster.

**Step 9** In the **Backup Data** column of the **Data Configuration** page, select one or more pieces of backup data that needs to be restored based on service requirements. In the **Target Path** column, specify the target location after backup data restoration.

You are advised to set **Target Path** to a new path that is different from the backup path.

**Step 10** Set **Force recovery** to **true**, which indicates that all backup data is forcibly restored when a data table with the same name already exists. If the data table contains new data added after backup, the new data will be lost after the data restoration. If you set the parameter to **false**, the restoration task is not executed if a data table with the same name exists.

**Step 11** Click **Verify** to check whether the restoration task is configured correctly.

- If the queue name is incorrect, the verification fails.
- If the specified directory to be restored does not exist, the verification fails.

**Step 12** Click **OK**.

**Step 13** In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

### 7.6.6.15 Restoring Kafka Metadata

#### Scenario

Kafka data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in ZooKeeper, an exception occurs or the operation has not achieved the expected result. All Kafka modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover Kafka data. Only manual restoration tasks are supported.

---

#### NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
  - To restore Kafka metadata when the service is running properly, you are advised to manually back up the latest Kafka metadata before restoration. Otherwise, the Kafka metadata that is generated after the data backup and before the data restoration will be lost.
-

## Impact on the System

- After the metadata is restored, the data generated after the data backup and before the data restoration is lost.
- After the metadata is restored, the offset information stored on ZooKeeper by Kafka consumers is rolled back, resulting in repeated consumption.

## Prerequisites

- If you need to restore data from a remote HDFS, a standby cluster has been created and the data has been backed up. For details, see [Backing Up Kafka Metadata](#). If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see [Configuring Mutual Trust Between MRS Clusters](#). If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see [Enabling MRS Inter-Cluster Replication](#).
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The Kafka service is disabled first, and then enabled upon data restoration.

## Restoring Kafka Metadata

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the **Operation** column of a specified task in the task list, choose **More > View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.  
Select the correct item, and manually copy the full path of backup files in **Backup Path**.

**Step 3** On FusionInsight Manager, choose **O&M > Backup and Restoration > Restoration Management**.

**Step 4** Click **Create**.

**Step 5** Set **Task Name** to the name of the restoration task.

**Step 6** Select the cluster to be operated from **Recovery Object**.

**Step 7** In the **Restoration Configuration** area, select **Kafka**.

**Step 8** Set **Path Type** of **Kafka** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.

If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Version\_Data source\_Task execution time.tar.gz*.

- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

If you select **LocalHDFS**, set the following parameters:

- **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/Backup task name\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **Source NameService Name**: indicates the NameService name that corresponds to the backup directory when a restoration task is executed. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select **RemoteHDFS**, set the following parameters:

- **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.

- **NFS**: indicates that backup files are stored in NAS using the NFS protocol.

If you select **NFS**, set the following parameters:

- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Server IP Address**: indicates the IP address of the NAS server.
- **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.

- **CIFS**: indicates that backup files are stored in NAS using the CIFS protocol.

If you select **CIFS**, set the following parameters:

- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.

- **Server IP Address:** indicates the IP address of the NAS server.
- **Port:** indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
- **Username:** indicates the username set when the CIFS protocol is configured.
- **Password:** indicates the password set when the CIFS protocol is configured.
- **Source Path:** indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.
- **OBS:** indicates that backup files are stored in OBS.  
If you select **OBS**, set the following parameters:
  - **Source Path:** indicates the full OBS path of a backup file, for example, *Backup path/Backup task name\_Data source\_Task creation time/Version\_Data source\_Task execution time.tar.gz*.

 **NOTE**

Only MRS 3.1.0 or later supports saving backup files in OBS.

**Step 9** Click **OK**.

**Step 10** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

**NOTICE**

- If the Kafka service is reinstalled and metadata is restored after data backup, or metadata is migrated to a new cluster, the Kafka broker cannot be restarted. View the error in the `/var/log/Bigdata/kafka/broker/server.log` file. An example is as follows:  
ERROR Fatal error during KafkaServer startup. Prepare to shutdown  
(kafka.server.KafkaServer)kafka.common.InconsistentClusterIdException: The Cluster ID kVSgfuUQFGGpHMTBqBPiw doesn't match stored clusterId Some(0Qftv9yBTAmf2iDPSllk7g) in meta.properties. The broker is trying to join the wrong cluster. Configured zookeeper.connect may be wrong. at kafka.server.KafkaServer.startup(KafkaServer.scala:220) at kafka.server.KafkaServerStartable.startup(KafkaServerStartable.scala:44) at kafka.Kafka\$.main(Kafka.scala:84) at kafka.Kafka.main(Kafka.scala)  
Check the value of `log.dirs` in the Kafka Broker configuration file `${BIGDATA_HOME}/FusionInsight_Current/*Broker/etc/server.properties`. The value is the Kafka data directory. Go to the Kafka data directory and change the value `0Qftv9yBTAmf2iDPSllk7g` of `cluster.id` in `meta.properties` to `kVSgfuUQFGGpHMTBqBPiw` (the latest value in the error log).
- The preceding modification must be performed on each node where Broker is located. After the modification, restart the Kafka service.

----End

## 7.6.7 Managing MRS Cluster Backup and Restoration Tasks

This topic describes how to modify the parameters of a created backup task on Manager to meet changing service requirements. The parameters of restoration tasks can be viewed but not modified.

### Impact on the System

After a backup task is modified, the new parameters take effect when the task is executed next time.

### Prerequisites

- A backup task has been created.
- A new backup task policy has been planned based on the actual situation.

### Modifying a Backup Task (MRS 3.x and Later)

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

**Step 2** In the task list, locate a specified task, click **Configure** in the **Operation** column to go to the configuration modification page.

On the displayed page, modify the following parameters:

- Started
- Period
- Destination NameService Name

- Target NameNode IP Address
- Target Path
- Max Number of Backup Copies
- Maximum Number of Recovery Points
- Maximum Number of Maps
- Maximum Bandwidth of a Map

 **NOTE**

After the **Target Path** parameter of a backup task is modified, this task will be performed as a full backup task for the first time by default.

**Step 3** Click **OK** to save the settings.

----End

## Modifying a Backup Task (MRS 2.x and Earlier)

**Step 1** On MRS Manager, choose **System > Back Up Data**.

**Step 2** In the task list, locate a specified task, click **Modify** in the **Operation** column to go to the configuration modification page.

**Step 3** Modify the following parameters on the displayed page:

- Manual backup:
  - Target Path
  - Max Number of Backup Copies
- Periodic backup:
  - Started
  - Period
  - Target Path
  - Max Number of Backup Copies

 **NOTE**

- When **Path Type** is set to **LocalHDFS**, **Target Path** is valid for modifying a backup task.
- After the **Target Path** parameter of a backup task is modified, this task will be performed as a full backup task for the first time by default.

**Step 4** Click **OK** to save the settings.

----End


## Viewing Backup and Restoration Tasks (MRS 3.x and Later)

**Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration**.

**Step 2** Click **Backup Management** or **Restoration Management**.

**Step 3** In the task list, obtain the previous execution result in the **Task Status** and **Task Progress** column. Green indicates that the task is executed successfully, and red indicates that the execution fails.

**Step 4** In the **Operation** column of a specified task in the task list, choose **More > View History** or click **View History** to view the historical record of backup and restoration task execution.

In the displayed window, click  before a specified record to display log information about the execution.

You can also perform more maintenance operations listed in [Table 7-34](#) in the **Operation** column of a specified task.

**Table 7-34** Maintenance and management operations

| Operation Entry                                      | Description                                                                                                                                                                                                                                                                              |
|------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>More &gt; Back Up Now</b> or <b>Start</b>         | Start a backup or restoration task that is ready or fails to be executed. Executed restoration tasks cannot be repeatedly executed.                                                                                                                                                      |
| <b>Configure</b>                                     | Modify parameters for a backup task.                                                                                                                                                                                                                                                     |
| <b>Restore</b>                                       | After data backup, click <b>Restore</b> to quickly restore data.                                                                                                                                                                                                                         |
| <b>More &gt; Back Up Now</b>                         | Execute a backup task immediately.                                                                                                                                                                                                                                                       |
| <b>More &gt; Stop</b> or <b>Stop</b>                 | Stop a running backup or restoration task. After the task is stopped, its <b>Task Status</b> changes to <b>Stopped</b> .                                                                                                                                                                 |
| <b>More &gt; Delete</b> or <b>Delete</b>             | Delete a backup or restoration task. After a task is deleted, the backup data is retained by default.                                                                                                                                                                                    |
| <b>More &gt; Suspend</b>                             | Suspend a backup task. Only periodic backup tasks can be suspended. Suspended backup tasks are no longer executed automatically. If you suspend a backup task that is being executed, the task execution stops. To unlock the task and run it again, choose <b>More &gt; Run Again</b> . |
| <b>More &gt; Resume</b>                              | Enable automatic backup.                                                                                                                                                                                                                                                                 |
| <b>More &gt; View History</b> or <b>View History</b> | Switch to the task run log page to view the task running details and backup path.                                                                                                                                                                                                        |
| <b>View</b>                                          | Check the parameter settings of a restoration task.                                                                                                                                                                                                                                      |
| <b>Start</b>                                         | Run a restoration task.                                                                                                                                                                                                                                                                  |

----End

## Viewing Backup and Restoration Tasks (MRS 2.x and Earlier)

**Step 1** On MRS Manager, click **System**.

**Step 2** Click **Backup Management** or **Restoration Management**.



**Step 3** In the task list, obtain the previous execution result in the **Task Progress** column. Green indicates that the task is executed successfully, and red indicates that the execution fails.

**Step 4** In the **Operation** column of a specified task in the task list, choose **More > View History** to view the historical record of backup and restoration execution.

In the displayed window, click **View** in the **Details** column. The task execution logs and paths are displayed.

#### NOTE

Other related operations are as follows:

- Viewing a restoration task

In the **Operation** column of the specified task in the task list, click **View Details** to view the restoration task. You can only view but cannot modify the parameters of a restoration task.

- Executing a backup or restoration task

In the task list, locate a specified task and click **Start** in the **Operation** column to start a backup or restoration task that is ready or fails to be executed. Executed restoration tasks cannot be repeatedly executed.

- Stopping a backup task

In the task list, locate a specified task and choose **More > Stop** in the **Operation** column to stop a backup task that is running.

- Deleting a backup or restoration task

In the **Operation** column of the specified task in the task list, choose **More > Delete** to delete the backup or restoration task. After a task is deleted, the backup data is retained by default.

- Suspending a backup task

In the task list, locate a specified task and choose **More > Suspend** in the **Operation** column to suspend the backup task. Only periodic backup tasks can be suspended. Suspended backup tasks are no longer executed automatically. When you suspend a backup task that is being executed, the task execution stops. To cancel the suspension status of a task, choose **More > Resume**.

----End

## 7.6.8 Using HDFS Snapshots to Quickly Restore Component Service Data

### Scenario

When DistCp is used to back up data, the backup snapshot is saved to HDFS of the active cluster. FusionInsight Manager supports using the local snapshot for quick data restoration, requiring less time than restoring data from the standby cluster.

Use FusionInsight Manager and the snapshots on HDFS of the active cluster to create a local quick restoration task and execute the task.

#### NOTE

This topic is available for MRS 3.x or later.

## Procedure

- Step 1** On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.
- Step 2** In the backup task list, locate a created task and click **Restore** in the **Operation** column.
- Step 3** Check whether the system displays "No data is available for quick restoration. Create a task on the restoration management page to restore data."
- If yes, click **OK** to close the dialog box. No backup data snapshot is created in the active cluster, and no further action is required.
  - If no, go to **Step 4** to create a local quick restoration task.

 **NOTE**

Metadata does not support quick restoration.

- Step 4** Set **Name** to the name of the local quick restoration task.
- Step 5** Set **Configuration** to a data source.
- Step 6** Set **Recovery Point List** to a recovery point that contains the backup data.
- Step 7** Set **Queue Name** to the name of the Yarn queue used in the task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- Step 8** Set **Data Configuration** to the object to be recovered.
- Step 9** Click **Verify**, and wait for the system to display "The restoration task configuration is verified successfully."
- Step 10** Click **OK**.
- Step 11** In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

After the task is complete, **Task Status** of the task is displayed as **Successful**.

----End

## 7.7 MRS Cluster Patching

### 7.7.1 Viewing Patch Information for an MRS Cluster

View patch information for cluster components. If a cluster component, such as Hadoop or Spark, is abnormal, download the patch version to rectify the fault.

- Step 1** Log in to the MRS management console.
- Step 2** On the **Active Clusters** page, click the name of your desired cluster.
- Step 3** Click **Patches**. On the displayed page, view patch information for the MRS cluster.
- Patch Name: name of the patch package
  - Published: time when the patch package is released

- Status: patch status
- Patch Description: patch version description
- Operation: patch installation or uninstallation

----End

## 7.7.2 Patching an MRS Cluster

Install the cluster version patch as required if you obtain patch information from:

- The message center service
- On a cluster details page of the MRS console, choose **Patches > Cluster Component Patches** to view the patches that can be installed.

### Preparing for Patch Installation

---

#### CAUTION

- For details about how to check the cluster status, see section [Performing a Health Check for an MRS Cluster](#). Exceptions such as cluster node faults and hard disk faults may cause patch installation and uninstallation failures. Before you install or uninstall the patch, ensure that the cluster is healthy.
  - Click **Patches** then **Cluster Component Patches**, view the **Patch Description** column of the target patch, and read the patch description carefully to understand the patch installation procedure and impact. For details, see [MRS Cluster Patch Description](#).
  - MRS 2.x and earlier versions, MRS 3.1.5 and later versions, and MRS 3.2.0-LTS and later versions support online patch installation.
- 

### Installing a Patch

- Step 1** Log in to the MRS management console.
- Step 2** On the **Active Clusters** page displayed by default, click the name of the target cluster to enter the cluster details page.
- Step 3** In the **Patches** tab, click **Cluster Component Patches**. In the operation list, click **Install** next to the patch you want to install.

#### NOTE

[MRS Cluster Patch Description](#) describes the patching procedure and impact.

- Step 4** In the displayed dialog box, select **I have read Patch Description and understood that this operation may restart services**. Click **Yes** and wait until the patch is successfully installed.
- Step 5** Check the patch status. Restart the components and install the client patch according to the patch description.

 NOTE

If there is an isolated host in the cluster, the patch will not be installed on the isolated host. In this case, when the installation completes, the patch is partially installed. After an isolated node is restored and the isolation is canceled, you can install the patch again. In this case, the patch is installed only on the node where isolation is canceled. For versions earlier than MRS 3.x, perform operations by referring to [Patching Hosts Isolated in an MRS Cluster](#).

----End

## Uninstalling a Patch

- Step 1** Log in to the MRS management console.
- Step 2** On the **Active Clusters** page displayed by default, click the name of the target cluster to enter the cluster details page.
- Step 3** In the **Patches** tab, click **Cluster Component Patches**. In the operation list, click **Uninstall** next to the patch you want to uninstall.
- Step 4** In the displayed dialog box, select the confirmation check box and click **Yes**. Wait until the patch is successfully uninstalled.
- Step 5** Restart the component and uninstall the client patch according to the patch description.

 NOTE

If there is an isolated host in the cluster, the patch will not be uninstalled on the isolated host. In this case, when the uninstallation completes, the patch is partially uninstalled. After the isolated node is restored and the isolation is canceled, you can uninstall the patch again. In this case, the patch is uninstalled only on the node whose isolation is canceled. For versions earlier than MRS 3.x, perform operations by referring to [Patching Hosts Isolated in an MRS Cluster](#).

----End

## 7.7.3 Applying Rolling Patches for an MRS Cluster

The rolling patch function indicates that patches are installed or uninstalled for one or more services in a cluster by performing a rolling service restart (restarting services or instances in batches), without interrupting the services or within a minimized service interruption interval. Services in a cluster are divided into the following three types based on whether they support rolling patch:

- Services supporting rolling patch installation or uninstallation: All businesses or part of them (varying depending on different services) of the services are not interrupted during patch installation or uninstallation.
- Services not supporting rolling patch installation or uninstallation: Businesses of the services are interrupted during patch installation or uninstallation.
- Services with some roles supporting rolling patch installation or uninstallation: Some businesses of the services are not interrupted during patch installation or uninstallation.

 NOTE

In **MRS 3.x**, you cannot perform operations in this section on the management console.

**Table 7-35** provides services and instances that support or do not support rolling restart in the MRS cluster.

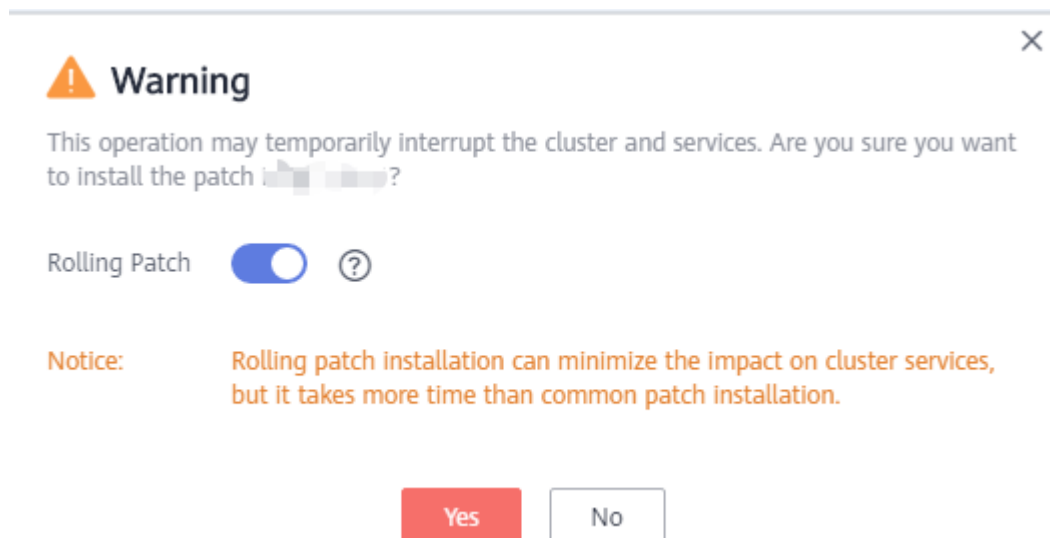
**Table 7-35** Support for rolling restarts on services and instances

| Service | Instance         | Support for Rolling Restart |
|---------|------------------|-----------------------------|
| Alluxio | AlluxioJobMaster | Yes                         |
|         | AlluxioMaster    |                             |
| Flink   | FlinkResource    | No                          |
|         | FlinkServer      |                             |
| Flume   | Flume            | Yes                         |
|         | MonitorServer    |                             |
| HBase   | HMaster          | Yes                         |
|         | RegionServer     |                             |
|         | ThriftServer     |                             |
|         | RETSerServer     |                             |
| HDFS    | NameNode         | Yes                         |
|         | Zkfc             |                             |
|         | JournalNode      |                             |
|         | HttpFS           |                             |
|         | DataNode         |                             |
| Hive    | MetaStore        | Yes                         |
|         | WebHCat          |                             |
|         | HiveServer       |                             |
| Hue     | Hue              | No                          |
| Impala  | Impalad          | No                          |
|         | StateStore       |                             |
|         | Catalog          |                             |
| Kafka   | Broker           | Yes                         |
|         | KafkaUI          | No                          |
| Kudu    | KuduTserver      | Yes                         |
|         | KuduMaster       |                             |
| Loader  | Sqoop            | No                          |

| Service   | Instance         | Support for Rolling Restart |
|-----------|------------------|-----------------------------|
| Mapreduce | JobHistoryServer | Yes                         |
| Oozie     | oozie            | No                          |
| Presto    | Coordinator      | Yes                         |
|           | Worker           |                             |
| Spark     | JobHistory       | Yes                         |
|           | JDBCServer       |                             |
|           | SparkResource    |                             |
| Storm     | Nimbus           | Yes                         |
|           | UI               |                             |
|           | Supervisor       |                             |
|           | Logviewer        |                             |
| Tez       | TezUI            | No                          |
| Yarn      | ResourceManager  | Yes                         |
|           | NodeManager      |                             |
| Zookeeper | Quorumpeer       | Yes                         |

## Installing a Rolling Patch

- Step 1** Log in to the MRS console.
- Step 2** On the **Active Clusters** page displayed by default, click the name of the target cluster to enter the cluster details page.
- Step 3** On the **Patches** page, click **Install** in the **Operation** column.
- Step 4** On the **Warning** page, enable or disable **Rolling Patch**.

**Figure 7-24** Rolling patch installation**NOTE**

- Enabling the rolling patch installation function: Services are not stopped before patch installation, and rolling service restart is performed after the patch installation. This minimizes the impact on cluster services but takes more time than common patch installation.
- Disabling the rolling patch uninstallation function: All services are stopped before patch uninstallation, and all services are restarted after the patch uninstallation. This temporarily interrupts the cluster and the services but takes less time than rolling patch uninstallation.
- The rolling patch installation function is not available in clusters with less than two Master nodes and three Core nodes.

**Step 5** Click **Yes** to install the target patch.

**Step 6** View the patch installation progress.

1. Access MRS Manager. For details, see [Accessing MRS Manager](#).
2. Choose **System > Manage Patch**. On the **Manage Patch** page, you can view the patch installation progress.

**NOTE**

For the isolated host nodes in the cluster, follow instructions in [Patching Hosts Isolated in an MRS Cluster](#) to restore the patch.

----End

## Uninstalling a Rolling Patch


**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page displayed by default, click the name of the target cluster to enter the cluster details page.

**Step 3** On the **Patches** page, click **Uninstall** in the **Operation** column.

**Step 4** On the **Warning** page, enable or disable **Rolling Patch**.

**Figure 7-25** Rolling patch uninstallation



**NOTE**

- Enabling the rolling patch uninstallation function: Services are not stopped before patch uninstallation, and rolling service restart is performed after the patch uninstallation. This minimizes the impact on cluster services but takes more time than common patch uninstallation.
- Disabling the rolling patch uninstallation function: All services are stopped before patch uninstallation, and all services are restarted after the patch uninstallation. This temporarily interrupts the cluster and the services but takes less time than rolling patch uninstallation.
- Only patches that are installed in rolling mode can be uninstalled in the same mode.

**Step 5** Click **Yes** to uninstall the target patch.

**Step 6** View the patch uninstallation progress.

1. Access MRS Manager. For details, see [Accessing MRS Manager](#).
2. Choose **System** > **Manage Patch**. On the **Manage Patch** page, you can view the patch uninstallation progress.

**NOTE**

For the isolated host nodes in the cluster, follow instructions in [Patching Hosts Isolated in an MRS Cluster](#) to restore the patch.

----End

## 7.7.4 Patching Hosts Isolated in an MRS Cluster

If some hosts are isolated in a cluster, perform the following operations to restore patches for these isolated hosts after patch installation on other hosts in the cluster. After patch restoration, versions of the isolated host nodes are consistent with those are not isolated.




**NOTE**

Operations in this section cannot be performed on the management console of MRS 3.x. This section applies only to versions earlier than 3.x.

- Step 1** Access MRS Manager. For details, see [Accessing MRS Manager](#).
- Step 2** Choose **System > Manage Patch**. The **Manage Patch** page is displayed.
- Step 3** In the **Operation** column, click **View Details**.
- Step 4** On the patch details page, select host nodes whose **Status** is **Isolated**.
- Step 5** Click **Select and Restore** to restore the isolated host nodes.

**Figure 7-26** Restoring patches for the isolated hosts



----End

## 7.8 MRS Cluster Patch Description

### 7.8.1 MRS 3.0.5.1 Patch Description

#### Basic Information

**Table 7-36** Basic information

|                      |             |
|----------------------|-------------|
| <b>Patch Version</b> | MRS 3.0.5.1 |
| <b>Release Date</b>  | 2021-08-14  |

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Resolved Issues</b></p> | <p><b>List of resolved issues in MRS 3.0.5.1:</b></p> <p><b>MRS Manager</b></p> <ul style="list-style-type: none"> <li>● Resolved the failure to submit SparkSQL jobs on the job management page due to long SQL statements.</li> <li>● Resolved the failure to execute SQL statements with comments.</li> </ul> <p><b>Big data components</b></p> <ul style="list-style-type: none"> <li>● Resolved the failure to synchronize IAM users with ClickHouse clusters.</li> <li>● Resolved the issue that the Flume client in the cluster cannot use an agency to access OBS.</li> <li>● Resolved the issue that the value of <b>% of Queue</b> is not displayed for a specified job on the native Yarn web UI.</li> <li>● Resolved the issue that job logs are incompletely displayed on the native Yarn web UI.</li> <li>● Resolved the issue that temporary files reside in HDFS after execution of Hive jobs.</li> <li>● Resolved the incompatibility in the interconnection between open-source Sqoop 1.4.7 and MRS Hive.</li> <li>● Resolved the failure to query Avro tables through Hive on MR.</li> <li>● Resolved the memory leak issue caused when HiveServer loads user-defined functions (UDFs).</li> <li>● Resolved the issue that the execution results of Hive and SparkSQL time functions are inconsistent.</li> <li>● Resolved the issue (HIVE-20187) that the result is incorrect when Hive on Tez uses MapJoin to achieve performance tuning.</li> <li>● Resolved the issue that an error occurs when the <b>beeline -p</b> command is executed.</li> <li>● Resolved the issue that Hue fails to format SQL statements.</li> <li>● Resolved the failure to submit Oozie jobs due to the incompatibility between Hue and Oozie time zones.</li> <li>● Resolved the unavailability of the variable drop-down list when a variable-declared Hive SQL statement is executed on the Hue web UI.</li> <li>● Resolved the query failure caused by incorrectly closed sessions when Hue connects to Hive for queries.</li> <li>● Resolved slow responses to Kunpeng servers' queries of Kudu tables using Impala.</li> <li>● Resolved the failure to install the Kudu client.</li> <li>● Resolved the unexpected restarts of KuduMaster instance on Kunpeng servers.</li> <li>● Resolved the search exceptions on the Ranger web UI.</li> <li>● Resolved the failure to redirect users to the login page after the logout from the Ranger web UI.</li> </ul> |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                                         |                                                                         |
|-----------------------------------------|-------------------------------------------------------------------------|
| <b>Compatibility with Other Patches</b> | The MRS 3.0.5.1 patch can resolve all the issues detected in MRS 3.0.5. |
|-----------------------------------------|-------------------------------------------------------------------------|

## Impact of Patch Installation

- During the installation of MRS 3.0.5.1, the executor and controller processes are automatically restarted and cluster functions on the management plane, such as job submission and cluster scaling, will be affected. Therefore, install the patch at an appropriate time.
- After you install the patch, restart the Spark2x, Hive, Yarn, Impala, Kudu, and Hue components on FusionInsight Manager for the patch to take effect. During the restart, some services may be unavailable for a short period of time. To minimize the impact on service continuity, perform the restart at a proper time.
- To install the MRS 3.0.5.1 patch, you need to manually download the patch file and install it on any master node in the cluster. For details, see the **README.md** file in the patch package.
- This patch must be also installed for any new node subsequently added to the cluster. To install the patch for this new node, install the patch on the master node and restart the corresponding service.

## Patch Download Addresses

- CN-Hong Kong: [https://mrs-container1-patch-ap-southeast-1.obs.ap-southeast-1.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_3.0.5.1\\_Patch\\_All\\_20210724.tar.gz](https://mrs-container1-patch-ap-southeast-1.obs.ap-southeast-1.myhuaweicloud.com/MRS_Common_Script/MRS_3.0.5.1_Patch_All_20210724.tar.gz)
- AP-Singapore: [https://mrs-container1-patch-ap-southeast-3.obs.ap-southeast-3.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_3.0.5.1\\_Patch\\_All\\_20210724.tar.gz](https://mrs-container1-patch-ap-southeast-3.obs.ap-southeast-3.myhuaweicloud.com/MRS_Common_Script/MRS_3.0.5.1_Patch_All_20210724.tar.gz)
- AP-Bangkok: [https://mrs-container1-patch-ap-southeast-2.obs.ap-southeast-2.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_3.0.5.1\\_Patch\\_All\\_20210724.tar.gz](https://mrs-container1-patch-ap-southeast-2.obs.ap-southeast-2.myhuaweicloud.com/MRS_Common_Script/MRS_3.0.5.1_Patch_All_20210724.tar.gz)

## 7.8.2 MRS 2.1.0.11 Patch Description

### Basic Information

Table 7-37 Basic information

|                      |              |
|----------------------|--------------|
| <b>Patch Version</b> | MRS 2.1.0.11 |
| <b>Release Date</b>  | 2020-12-30   |

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Resolved Issues</b></p> | <p><b>List of resolved issues in MRS 2.1.0.11:</b></p> <p><b>MRS Manager</b></p> <p>Executor, KNOX, and OS logs can be rolled back.<br/>           Executor GC logs are now added.<br/>           The Knox restart failure is resolved.<br/>           Resolve the problem that jobs fail to be submitted when a node is faulty.<br/>           Full-link monitoring is supported.<br/>           Job Status can be updated when switching over the ResourceManager active and standby nodes.<br/>           Backup and restoration fail in some scenarios is resolved.<br/>           The process fault alarm that frequently generated on the HMaster is resolved.</p> <p><b>Big data components</b></p> <p>Resolved the issue of JobHistory memory leakage.<br/>           Resolved the issue of the Hive truncate table times out and fails to be truncated.<br/>           Resolved the issue that the table data file does not exist after an incremental Hive task fails.<br/>           The Hive SQL statement is not running properly.<br/>           After a Carbon table is created in a security cluster and the hive group does not have the permission to create the Carbon table, other users can create the Carbon table.<br/>           Resolved the problem that the spark JDBCServer process is abnormal.</p> |
|                               | <p><b>List of resolved issues in MRS 2.1.0.10:</b></p> <p><b>MRS Manager</b></p> <p>New queue configurations in the <b>capacity-schedule.xml</b> file will not be lost during cluster scale-out after the patch is installed.<br/>           Full-link monitoring can be rolled back.</p> <p><b>Big data components</b></p> <p>Hive permission assignment failure on Spark is resolved.<br/>           If no queue is specified, tasks are submitted to the launcher-job queue by default. Task running will not be affected.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 2.1.0.9:</b></p> <p><b>MRS Manager</b></p> <p>The MRS Executor memory overflow is resolved.</p> <p>The cluster scale-out process is optimized.</p> <p>The problem that the SQL statement is incorrectly combined when the value of SparkSQL contains spaces is resolved.</p> <p>The problem that HiveSQL jobs fail to be submitted occasionally is resolved.</p> <p>The permission control for downloading the keytab file is optimized.</p> <p><b>Big data components</b></p> <p>When the Presto role name contains uppercase letters, the permission model can take effect.</p> <p>The problem that Hive partitions are deleted slowly is resolved.</p> <p>The problem that the token expires after Spark runs for a long time is resolved.</p> |
|  | <p><b>List of resolved issues in MRS 2.1.0.8:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the ECS API traffic is limited when OBS is accessed through an agency has been solved.</p> <p>Multiple users can log in to MRS Manager at the same time.</p> <p>Full-link monitoring is supported.</p> <p><b>MRS big data components</b></p> <p>Carbon 2.0 has been upgraded.</p> <p>The HBASE-18484 issue has been solved.</p>                                                                                                                                                                                                                                                                                                                                                    |
|  | <p><b>List of the resolved issues in MRS 2.1.0.7:</b></p> <p><b>MRS Manager</b></p> <p>The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved.</p> <p>The Presto query result can be saved as a file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 2.1.0.6:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the disk I/O usage of monitoring data is inaccurate has been solved.</p> <p>The problem that the Spark job status is not updated occasionally has been solved.</p> <p>The problem that the job running failure has been solved.</p> <p>The patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The HBase exceptions are rectified.</p> <p>The problem that the system responds slowly when Hive roles are bound to permissions has been solved.</p> |
|  | <p><b>List of resolved issues in MRS 2.1.0.5:</b></p> <p><b>MRS big data components</b></p> <p>Impala supports the ObsFileSystem function.</p> <p>The timeout period of the MRS Manager page and the native pages of components can be configured.</p> <p>The Hive privilege binding freezing problem has been solved.</p> <p>The data connection failure has been solved.</p>                                                                                                                                                                                                       |
|  | <p><b>List of resolved issues in MRS 2.1.0.3:</b></p> <p><b>MRS Manager</b></p> <p>Problems of Manager executor's high concurrent job submission have been solved.</p> <p><b>MRS big data components</b></p> <p>Data insertion failure in hive on tez has been fixed.</p>                                                                                                                                                                                                                                                                                                            |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.2:</b></p> <p><b>MRS Manager</b></p> <p>No monitoring information is displayed after NodeAgent is restarted.</p> <p>When a job is under submission for a long time, memory overflow occurs in the <b>manager executor</b> process.</p> <p>Job submission is supported. <b>manager executor</b> can be used to configure high concurrency.</p> <p>New Kafka topics are not displayed on the MRS Manager management plane.</p> <p>When you call security cluster's APIs to submit the <b>Spark Submit</b> job and perform operations on an HBase table, the permission control on the HBase table does not take effect.</p> <p>The MRS Manager patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The slow running of the <b>load data inpath</b> command executed by Spark has been optimized.</p> <p>Column names containing the dollar sign (\$) can be used in Spark table creation.</p> <p>OBS-related problems have been solved.</p> |
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b></p> <p>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b></p> <p>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 2.1.0.11 patch package contains all patches released for MRS 2.1.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <p><b>Vulnerab<br/>ility<br/>Disclosur<br/>e</b></p>            | <p>The remote code execution vulnerability of the Spark has been fixed. For details about the vulnerability, see <a href="#">CVE-2020-9480</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## Impact of Patch Installation

- During the installation of the MRS 2.1.0.11 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, YARN, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are

temporarily unavailable but services are not interrupted during the rolling restart.

- After the MRS 2.1.0.11 patch is installed, log in to the **standby Master node** (Log in to MRS Manager. The Master node with a hollow pentagon on the **Host Management** page is the standby Master node), switch to user **omm**, and run the **sh /opt/knox/bin/restart-knox.sh** command to restart the Knox process. This operation is not required for a cluster with only one Master node.

You can run the **ps -ef |grep knox** command to check whether the knox process is started. If the knox process ID is displayed, the knox process is started successfully.

- (Optional) After installing the MRS 2.1.0.11 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

#### NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:
  - a. Change the session timeout interval of the web and CAS services on all Master nodes.
    - i. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml`. The unit is minute.
    - ii. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml`. The unit is minute.
  - b. Change the TGT validity period of the CAS on all Master nodes.  
Change `1200` in `p:maxTimeToLiveInSeconds="$ {tgt.maxTimeToLiveInSeconds:1200}` and `p:timeToKillInSeconds="$ {tgt.timeToKillInSeconds:1200}"` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml` to the corresponding timeout interval, in seconds.
  - c. Restart the Tomcat service on the active Master node.
    - i. On the active Master node, run the `netstat -anp |grep 28443 |grep LISTEN` command as user **omm** to query the Tomcat process ID.
    - ii. Run the `kill -9 {pid}` command, in which `{pid}` indicates the process ID obtained in the previous step.



- iii. Wait for the process to automatically restart. You can run the **netstat -anp |grep 28443 |grep LISTEN** command to check whether the process is started. If the command output is displayed, the process is started successfully.
- d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.
  - HDFS/MapReduce/YARN: Add the custom configuration item **http.server.session.timeout.secs**.
  - Spark: Change the value of **spark.session.maxAge**.
  - Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

## 7.8.3 MRS 2.1.0.10 Patch Description

### Basic Information

Table 7-38 Basic information

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Patch Version   | MRS 2.1.0.10                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Release Date    | 2020-09-21                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Resolved Issues | <b>List of resolved issues in MRS 2.1.0.10:</b><br><b>MRS Manager</b><br>New queue configurations in the <b>capacity-schedule.xml</b> file will not be lost during cluster scale-out after the patch is installed.<br>Full-link monitoring can be rolled back.<br><b>Big data components</b><br>Hive permission assignment failure on Spark is resolved.<br>If no queue is specified, tasks are submitted to the launcher-job queue by default. Task running will not be affected. |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 2.1.0.9:</b></p> <p><b>MRS Manager</b></p> <p>The MRS Executor memory overflow is resolved.</p> <p>Optimized the cluster scale-out process.</p> <p>The problem that the SQL statement is incorrectly combined when the value of SparkSQL contains spaces is resolved.</p> <p>The problem that HiveSQL jobs fail to be submitted occasionally is resolved.</p> <p>The permission control for downloading the keytab file is optimized.</p> <p><b>Big data components</b></p> <p>When the Presto role name contains uppercase letters, the permission model can take effect.</p> <p>The problem that Hive partitions are deleted slowly is resolved.</p> <p>The problem that the token expires after Spark runs for a long time is resolved.</p> |
|  | <p><b>List of resolved issues in MRS 2.1.0.8:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the ECS API traffic is limited when OBS is accessed through an agency has been solved.</p> <p>Multiple users can log in to MRS Manager at the same time.</p> <p>Full-link monitoring is supported.</p> <p><b>MRS big data components</b></p> <p>Carbon 2.0 has been upgraded.</p> <p>The HBASE-18484 issue has been solved.</p>                                                                                                                                                                                                                                                                                                                                                 |
|  | <p><b>List of the resolved issues in MRS 2.1.0.7:</b></p> <p><b>MRS Manager</b></p> <p>The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved.</p> <p>The Presto query result can be saved as a file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 2.1.0.6:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the disk I/O usage of monitoring data is inaccurate has been solved.</p> <p>The problem that the Spark job status is not updated occasionally has been solved.</p> <p>The problem that the job running failure has been solved.</p> <p>The patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The HBase exceptions are rectified.</p> <p>The problem that the system responds slowly when Hive roles are bound to permissions has been solved.</p> |
|  | <p><b>List of resolved issues in MRS 2.1.0.5:</b></p> <p><b>MRS big data components</b></p> <p>Impala supports the ObsFileSystem function.</p> <p>The timeout period of the MRS Manager page and the native pages of components can be configured.</p> <p>The Hive privilege binding freezing problem has been solved.</p> <p>The data connection failure has been solved.</p>                                                                                                                                                                                                       |
|  | <p><b>List of resolved issues in MRS 2.1.0.3:</b></p> <p><b>MRS Manager</b></p> <p>Problems of Manager executor's high concurrent job submission have been solved.</p> <p><b>MRS big data components</b></p> <p>Data insertion failure in <b>hive on tez</b> has been fixed.</p>                                                                                                                                                                                                                                                                                                     |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.2:</b></p> <p><b>MRS Manager</b></p> <p>No monitoring information is displayed after NodeAgent is restarted.</p> <p>When a job is under submission for a long time, memory overflow occurs in the <b>manager executor</b> process.</p> <p>Job submission is supported. <b>manager executor</b> can be used to configure high concurrency.</p> <p>New Kafka topics are not displayed on the MRS Manager management plane.</p> <p>When you call security cluster's APIs to submit the <b>Spark Submit</b> job and perform operations on an HBase table, the permission control on the HBase table does not take effect.</p> <p>The MRS Manager patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The slow running of the <b>load data inpath</b> command executed by Spark has been optimized.</p> <p>Column names containing the dollar sign (\$) can be used in Spark table creation.</p> <p>OBS-related problems have been solved.</p> |
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b></p> <p>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b></p> <p>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 2.1.0.10 patch package contains all patches released for MRS 2.1.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <p><b>Vulnerab<br/>ility<br/>Disclosur<br/>e</b></p>            | <p>The remote code execution vulnerability of the Spark has been fixed. For details about the vulnerability, see <a href="#">CVE-2020-9480</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## Impact of Patch Installation

- During the installation of the MRS 2.1.0.10 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, YARN, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are

temporarily unavailable but services are not interrupted during the rolling restart.

- After installing the MRS 2.1.0.10 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

#### NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:
  - a. Change the session timeout interval of the web and CAS services on all Master nodes.
    - i. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml`. The unit is minute.
    - ii. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml`. The unit is minute.
  - b. Change the TGT validity period of the CAS on all Master nodes.

Change `1200` in `p:maxTimeToLiveInSeconds="$ {tgt.maxTimeToLiveInSeconds:1200}` and `p:timeToKillInSeconds="$ {tgt.timeToKillInSeconds:1200}"` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml` to the corresponding timeout interval, in seconds.
  - c. Restart the Tomcat service on the active Master node.
    - i. On the active Master node, run the `netstat -anp |grep 28443 |grep LISTEN` command as user `omm` to query the Tomcat process ID.
    - ii. Run the `kill -9 {pid}` command, in which `{pid}` indicates the process ID obtained in the previous step.
    - iii. Wait for the process to automatically restart. You can run the `netstat -anp |grep 28443 |grep LISTEN` command to check whether the process is started. If the command output is displayed, the process is started successfully.
  - d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.
    - HDFS/MapReduce/YARN: Add the custom configuration item `http.server.session.timeout.secs`.

- Spark: Change the value of **spark.session.maxAge**.
- Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

## 7.8.4 MRS 2.1.0.9 Patch Description

### Basic Information

Table 7-39 Basic information

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>   | MRS 2.1.0.9                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Date</b>    | 2020-08-21                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Resolved Issues</b> | <p><b>List of resolved issues in MRS 2.1.0.9:</b></p> <p><b>MRS Manager</b></p> <p>The MRS Executor memory overflow is resolved.</p> <p>Optimized the disk scale-out process.</p> <p>The problem that the SQL statement is incorrectly combined when the value of SparkSQL contains spaces is resolved.</p> <p>The problem that HiveSQL jobs fail to be submitted occasionally is resolved.</p> <p>The permission control for downloading the keytab file is optimized.</p> <p><b>Big data components</b></p> <p>When the Presto role name contains uppercase letters, the permission model can take effect.</p> <p>Partitions are deleted slowly in Hive.</p> <p>The problem that the token expires after Spark runs for a long time is resolved.</p> <hr/> <p><b>List of resolved issues in MRS 2.1.0.8:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the ECS API traffic is limited when OBS is accessed through an agency has been solved.</p> <p>Multiple users can log in to MRS Manager at the same time.</p> <p>Full-link monitoring is supported.</p> <p><b>MRS big data components</b></p> <p>Carbon 2.0 has been upgraded.</p> <p>The HBASE-18484 issue has been solved.</p> |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of the resolved issues in MRS 2.1.0.7:</b></p> <p><b>MRS Manager</b></p> <p>The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved.</p> <p>The Presto query result can be saved as a file.</p>                                                                                                                                                                                                                                                                                         |
|  | <p><b>List of resolved issues in MRS 2.1.0.6:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the disk I/O usage of monitoring data is inaccurate has been solved.</p> <p>The problem that the Spark job status is not updated occasionally has been solved.</p> <p>The problem that the job running failure has been solved.</p> <p>The patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The HBase exceptions are rectified.</p> <p>The problem that the system responds slowly when Hive roles are bound to permissions has been solved.</p> |
|  | <p><b>List of resolved issues in MRS 2.1.0.5:</b></p> <p><b>MRS big data components</b></p> <p>Impala supports the ObsFileSystem function.</p> <p>The timeout period of the MRS Manager page and the native pages of components can be configured.</p> <p>The Hive privilege binding freezing problem has been solved.</p> <p>The data connection failure has been solved.</p>                                                                                                                                                                                                       |
|  | <p><b>List of resolved issues in MRS 2.1.0.3:</b></p> <p><b>MRS Manager</b></p> <p>Problems of Manager executor's high concurrent job submission have been solved.</p> <p><b>MRS big data components</b></p> <p>Data insertion failure in <b>hive on tez</b> has been fixed.</p>                                                                                                                                                                                                                                                                                                     |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.2:</b></p> <p><b>MRS Manager</b></p> <p>No monitoring information is displayed after NodeAgent is restarted.</p> <p>When a job is under submission for a long time, memory overflow occurs in the <b>manager executor</b> process.</p> <p>Job submission is supported. <b>manager executor</b> can be used to configure high concurrency.</p> <p>New Kafka topics are not displayed on the MRS Manager management plane.</p> <p>When you call security cluster's APIs to submit the <b>Spark Submit</b> job and perform operations on an HBase table, the permission control on the HBase table does not take effect.</p> <p>The MRS Manager patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The slow running of the <b>load data inpath</b> command executed by Spark has been optimized.</p> <p>Column names containing the dollar sign (\$) can be used in Spark table creation.</p> <p>OBS-related problems have been solved.</p> |
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b></p> <p>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b></p> <p>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 2.1.0.9 patch package contains all patches released for MRS 2.1.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <p><b>Vulnerab<br/>ility<br/>Disclosur<br/>e</b></p>            | <p>The remote code execution vulnerability of the Spark has been fixed. For details about the vulnerability, see <a href="#">CVE-2020-9480</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## Impact of Patch Installation

- During the installation of the MRS 2.1.0.9 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, YARN, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are



temporarily unavailable but services are not interrupted during the rolling restart.

- After installing the MRS 2.1.0.9 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

#### NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:
  - a. Change the session timeout interval of the web and CAS services on all Master nodes.
    - i. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml`. The unit is minute.
    - ii. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml`. The unit is minute.
  - b. Change the TGT validity period of the CAS on all Master nodes.

Change `1200` in `p:maxTimeToLiveInSeconds="$ {tgt.maxTimeToLiveInSeconds:1200}` and `p:timeToKillInSeconds="$ {tgt.timeToKillInSeconds:1200}`" in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml` to the corresponding timeout interval, in seconds.
  - c. Restart the Tomcat service on the active Master node.
    - i. On the active Master node, run the `netstat -anp |grep 28443 |grep LISTEN` command as user `omm` to query the Tomcat process ID.
    - ii. Run the `kill -9 {pid}` command, in which `{pid}` indicates the process ID obtained in the previous step.
    - iii. Wait for the process to automatically restart. You can run the `netstat -anp |grep 28443 |grep LISTEN` command to check whether the process is started. If the command output is displayed, the process is started successfully.
  - d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.
    - HDFS/MapReduce/YARN: Add the custom configuration item `http.server.session.timeout.secs`.

- Spark: Change the value of **spark.session.maxAge**.
- Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

## 7.8.5 MRS 2.1.0.8 Patch Description

### Basic Information

Table 7-40 Basic information

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>   | MRS 2.1.0.8                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Date</b>    | 2020-08-04                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Resolved Issues</b> | <p><b>List of resolved issues in MRS 2.1.0.8:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the ECS API traffic is limited when OBS is accessed through an agency has been solved.</p> <p>Multiple users can log in to MRS Manager at the same time.</p> <p>Full-link monitoring is supported.</p> <p><b>MRS big data components</b></p> <p>Carbon 2.0 has been upgraded.</p> <p>The HBASE-18484 issue has been solved.</p> <p><b>List of the resolved issues in MRS 2.1.0.7:</b></p> <p><b>MRS Manager</b></p> <p>The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved.</p> <p>The Presto query result can be saved as a file.</p> |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 2.1.0.6:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the disk I/O usage of monitoring data is inaccurate has been solved.</p> <p>The problem that the Spark job status is not updated occasionally has been solved.</p> <p>The problem that the job running failure has been solved.</p> <p>The patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The HBase exceptions are rectified.</p> <p>The problem that the system responds slowly when Hive roles are bound to permissions has been solved.</p> |
|  | <p><b>List of resolved issues in MRS 2.1.0.5:</b></p> <p><b>MRS big data components</b></p> <p>Impala supports the ObsFileSystem function.</p> <p>The timeout period of the MRS Manager page and the native pages of components can be configured.</p> <p>The Hive privilege binding freezing problem has been solved.</p> <p>The data connection failure has been solved.</p>                                                                                                                                                                                                       |
|  | <p><b>List of resolved issues in MRS 2.1.0.3:</b></p> <p><b>MRS Manager</b></p> <p>Problems of Manager executor's high concurrent job submission have been solved.</p> <p><b>MRS big data components</b></p> <p>Data insertion failure in hive on tez has been fixed.</p>                                                                                                                                                                                                                                                                                                            |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.2:</b></p> <p><b>MRS Manager</b></p> <p>No monitoring information is displayed after NodeAgent is restarted.</p> <p>When a job is under submission for a long time, memory overflow occurs in the <b>manager executor</b> process.</p> <p>Job submission is supported. <b>manager executor</b> can be used to configure high concurrency.</p> <p>New Kafka topics are not displayed on the MRS Manager management plane.</p> <p>When you call security cluster's APIs to submit the <b>Spark Submit</b> job and perform operations on an HBase table, the permission control on the HBase table does not take effect.</p> <p>The MRS Manager patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The slow running of the <b>load data inpath</b> command executed by Spark has been optimized.</p> <p>Column names containing the dollar sign (\$) can be used in Spark table creation.</p> <p>OBS-related problems have been solved.</p> |
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b></p> <p>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b></p> <p>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 2.1.0.8 patch package contains all patches released for MRS 2.1.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <p><b>Vulnerab<br/>ility<br/>Disclosur<br/>e</b></p>            | <p>The remote code execution vulnerability of the Spark has been fixed. For details about the vulnerability, see <a href="#">CVE-2020-9480</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## Impact of Patch Installation

- During the installation of the MRS 2.1.0.8 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, YARN, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are

temporarily unavailable but services are not interrupted during the rolling restart.

- After installing the MRS 2.1.0.8 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

#### NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:
  - a. Change the session timeout interval of the web and CAS services on all Master nodes.
    - i. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml`. The unit is minute.
    - ii. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml`. The unit is minute.
  - b. Change the TGT validity period of the CAS on all Master nodes.  
Change `1200` in `p:maxTimeToLiveInSeconds="$ {tgt.maxTimeToLiveInSeconds:1200}` and `p:timeToKillInSeconds="$ {tgt.timeToKillInSeconds:1200}"` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml` to the corresponding timeout interval, in seconds.
  - c. Restart the Tomcat service on the active Master node.
    - i. On the active Master node, run the `netstat -anp |grep 28443 |grep LISTEN` command as user `omm` to query the Tomcat process ID.
    - ii. Run the `kill -9 {pid}` command, in which `{pid}` indicates the process ID obtained in the previous step.
    - iii. Wait for the process to automatically restart. You can run the `netstat -anp |grep 28443 |grep LISTEN` command to check whether the process is started. If the command output is displayed, the process is started successfully.
  - d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.
    - HDFS/MapReduce/YARN: Add the custom configuration item `http.server.session.timeout.secs`.

- Spark: Change the value of **spark.session.maxAge**.
- Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

## 7.8.6 MRS 2.1.0.7 Patch Description

### Basic Information

Table 7-41 Basic information

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>   | MRS 2.1.0.7                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Date</b>    | 2020-07-15                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Resolved Issues</b> | <p><b>List of the resolved issues in MRS 2.1.0.7:</b></p> <p><b>MRS Manager</b></p> <p>The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved.</p> <p>The Presto query result can be saved as a file.</p>                                                                                                                                                                                                                                                                                         |
|                        | <p><b>List of resolved issues in MRS 2.1.0.6:</b></p> <p><b>MRS Manager</b></p> <p>The problem that the disk I/O usage of monitoring data is inaccurate has been solved.</p> <p>The problem that the Spark job status is not updated occasionally has been solved.</p> <p>The problem that the job running failure has been solved.</p> <p>The patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The HBase exceptions are rectified.</p> <p>The problem that the system responds slowly when Hive roles are bound to permissions has been solved.</p> |
|                        | <p><b>List of resolved issues in MRS 2.1.0.5:</b></p> <p><b>MRS big data components</b></p> <p>Impala supports the ObsFileSystem function.</p> <p>The timeout period of the MRS Manager page and the native pages of components can be configured.</p> <p>The Hive privilege binding freezing problem has been solved.</p> <p>The data connection failure has been solved.</p>                                                                                                                                                                                                       |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of resolved issues in MRS 2.1.0.3:</b></p> <p><b>MRS Manager</b><br/>Problems of Manager executor's high concurrent job submission have been solved.</p> <p><b>MRS big data components</b><br/>Data insertion failure in hive on tez has been fixed.</p> <hr/> <p><b>List of the resolved issues in MRS 2.1.0.2:</b></p> <p><b>MRS Manager</b><br/>No monitoring information is displayed after NodeAgent is restarted.<br/>When a job is under submission for a long time, memory overflow occurs in the <b>manager executor</b> process.<br/>Job submission is supported. <b>manager executor</b> can be used to configure high concurrency.<br/>New Kafka topics are not displayed on the MRS Manager management plane.<br/>When you call security cluster's APIs to submit the <b>Spark Submit</b> job and perform operations on an HBase table, the permission control on the HBase table does not take effect.<br/>The MRS Manager patch mechanism has been optimized.</p> <p><b>MRS big data components</b><br/>The slow running of the <b>load data inpath</b> command executed by Spark has been optimized.<br/>Column names containing the dollar sign (\$) can be used in Spark table creation.<br/>OBS-related problems have been solved.</p> <hr/> <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b><br/>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b><br/>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p> |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 2.1.0.7 patch package contains all patches released for MRS 2.1.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Impact of Patch Installation

- During the installation of the MRS 2.1.0.7 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, YARN,

MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.

- After installing the MRS 2.1.0.7 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

#### NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:
  - a. Change the session timeout interval of the web and CAS services on all Master nodes.
    - i. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml`. The unit is minute.
    - ii. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml`. The unit is minute.
  - b. Change the TGT validity period of the CAS on all Master nodes.  
Change `1200` in `p:maxTimeToLiveInSeconds="$ {tgt.maxTimeToLiveInSeconds:1200}` and `p:timeToKillInSeconds="$ {tgt.timeToKillInSeconds:1200}"` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-ticketExpirationPolicies.xml` to the corresponding timeout interval, in seconds.
  - c. Restart the Tomcat service on the active Master node.
    - i. On the active Master node, run the `netstat -anp |grep 28443 |grep LISTEN` command as user `omm` to query the Tomcat process ID.
    - ii. Run the `kill -9 {pid}` command, in which `{pid}` indicates the process ID obtained in the previous step.
    - iii. Wait for the process to automatically restart. You can run the `netstat -anp |grep 28443 |grep LISTEN` command to check whether the process is started. If the command output is displayed, the process is started successfully.
  - d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.
    - HDFS/MapReduce/YARN: Add the custom configuration item `http.server.session.timeout.secs`.



- Spark: Change the value of **spark.session.maxAge**.
- Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

## 7.8.7 MRS 2.1.0.6 Patch Description

### Basic Information

Table 7-42 Basic information

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>   | MRS 2.1.0.6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Date</b>    | 2020-06-10                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Resolved Issues</b> | <p><b>List of resolved issues in MRS 2.1.0.6:</b></p> <p><b>MRS Manager</b><br/>The problem that the disk I/O usage of monitoring data is inaccurate has been solved.<br/>The problem that the Spark job status is not updated occasionally has been solved.<br/>The problem that the job running failure has been solved.<br/>The patch mechanism has been optimized.</p> <p><b>MRS big data components</b><br/>The HBase exceptions are rectified.<br/>The problem that the system responds slowly when Hive roles are bound to permissions has been solved.</p> <hr/> <p><b>List of resolved issues in MRS 2.1.0.5:</b></p> <p><b>MRS big data components</b><br/>Impala supports the ObsFileSystem function.<br/>The timeout period of the MRS Manager page and the native pages of components can be configured.<br/>The Hive privilege binding freezing problem has been solved.<br/>The data connection failure has been solved.</p> <hr/> <p><b>List of resolved issues in MRS 2.1.0.3:</b></p> <p><b>MRS Manager</b><br/>Problems of Manager executor's high concurrent job submission have been solved.</p> <p><b>MRS big data components</b><br/>Data insertion failure in hive on tez has been fixed.</p> |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.2:</b></p> <p><b>MRS Manager</b></p> <p>No monitoring information is displayed after NodeAgent is restarted.</p> <p>When a job is under submission for a long time, memory overflow occurs in the <b>manager executor</b> process.</p> <p>Job submission is supported. <b>manager executor</b> can be used to configure high concurrency.</p> <p>New Kafka topics are not displayed on the MRS Manager management plane.</p> <p>When you call security cluster's APIs to submit the <b>Spark Submit</b> job and perform operations on an HBase table, the permission control on the HBase table does not take effect.</p> <p>The MRS Manager patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The slow running of the <b>load data inpath</b> command executed by Spark has been optimized.</p> <p>Column names containing the dollar sign (\$) can be used in Spark table creation.</p> <p>OBS-related problems have been solved.</p> |
|                                                                 | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b></p> <p>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b></p> <p>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 2.1.0.6 patch package contains all patches released for MRS 2.1.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Impact of Patch Installation

- During the installation of the MRS 2.1.0.6 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, YARN, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 2.1.0.6 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).

- For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
- For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

**NOTE**

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:
  - a. Change the session timeout interval of the web and CAS services on all Master nodes.
    - i. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml`. The unit is minute.
    - ii. Change the value of `<session-timeout>20</session-timeout>` in `/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml`. The unit is minute.
  - b. Change the TGT validity period of the CAS on all Master nodes.  
Change `1200` in `p:maxTimeToLiveInSeconds="$ {tgt.maxTimeToLiveInSeconds:1200}` and `p:timeToKillInSeconds="$ {tgt.timeToKillInSeconds:1200}"` in `/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml` to the corresponding timeout interval, in seconds.
  - c. Restart the Tomcat service on the active Master node.
    - i. On the active Master node, run the `netstat -anp |grep 28443 |grep LISTEN` command as user `omm` to query the Tomcat process ID.
    - ii. Run the `kill -9 {pid}` command, in which `{pid}` indicates the process ID obtained in the previous step.
    - iii. Wait for the process to automatically restart. You can run the `netstat -anp |grep 28443 |grep LISTEN` command to check whether the process is started. If the command output is displayed, the process is started successfully.
  - d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.
    - HDFS/MapReduce/YARN: Add the custom configuration item `http.server.session.timeout.secs`.
    - Spark: Change the value of `spark.session.maxAge`.
    - Hive: Add the customized configuration item `http.server.session.timeout.secs`.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

## 7.8.8 MRS 2.1.0.3 Patch Description

### Basic Information

Table 7-43 Basic information

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>   | MRS 2.1.0.3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Date</b>    | 2020-04-29                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Resolved Issues</b> | <p><b>List of resolved issues in MRS 2.1.0.3:</b></p> <p><b>MRS Manager</b><br/>Problems of Manager executor's high concurrent job submission have been solved.</p> <p><b>MRS big data components</b><br/>Data insertion failure in hive on tez has been fixed.</p> <hr/> <p><b>List of the resolved issues in MRS 2.1.0.2:</b></p> <p><b>MRS Manager</b><br/>No monitoring information is displayed after NodeAgent is restarted.<br/>When a job is under submission for a long time, memory overflow occurs in the <b>manager executor</b> process.<br/>Job submission is supported. <b>manager executor</b> can be used to configure high concurrency.<br/>New Kafka topics are not displayed on the MRS Manager management plane.<br/>When you call security cluster's APIs to submit the <b>Spark Submit</b> job and perform operations on an HBase table, the permission control on the HBase table does not take effect.<br/>The MRS Manager patch mechanism has been optimized.</p> <p><b>MRS big data components</b><br/>The slow running of the <b>load data inpath</b> command executed by Spark has been optimized.<br/>Column names containing the dollar sign (\$) can be used in Spark table creation.<br/>OBS-related problems have been solved.</p> |

|                                                |                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b></p> <p>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b></p> <p>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p> |
| <p><b>Compatibility with Other Patches</b></p> | <p>The MRS 2.1.0.3 patch package contains all patches released for MRS 2.1.0.</p>                                                                                                                                                                                                                                                                                                                          |

## Impact of Patch Installation

- During the installation of the MRS 2.1.0.3 patch, MRS Manager will be restarted, and the components such as Hive, Spark, HDFS, YARN, MapReduce, Presto, HBase, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 2.1.0.3 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

### NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

## 7.8.9 MRS 2.1.0.2 Patch Description

### Basic Information

**Table 7-44** Basic information

|                             |                    |
|-----------------------------|--------------------|
| <p><b>Patch Version</b></p> | <p>MRS 2.1.0.2</p> |
| <p><b>Release Date</b></p>  | <p>2020-04-22</p>  |

|                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Resolved Issues</b></p>                  | <p><b>List of the resolved issues in MRS 2.1.0.2:</b></p> <p><b>MRS Manager</b></p> <p>No monitoring information is displayed after NodeAgent is restarted.</p> <p>When a job is under submission for a long time, memory overflow occurs in the <b>manager executor</b> process.</p> <p>Job submission is supported. You can configure the concurrency for <b>manager executor</b>.</p> <p>New Kafka topics are not displayed on the MRS Manager management plane.</p> <p>When you call security cluster's APIs to submit the <b>Spark Submit</b> job and perform operations on an HBase table, the permission control on the HBase table does not take effect.</p> <p>The MRS Manager patch mechanism has been optimized.</p> <p><b>MRS big data components</b></p> <p>The slow running of the <b>load data inpath</b> command executed by Spark has been optimized.</p> <p>Column names containing the dollar sign (\$) can be used in Spark table creation.</p> <p>OBS-related problems have been solved.</p> |
|                                                | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b></p> <p>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b></p> <p>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <p><b>Compatibility with Other Patches</b></p> | <p>The MRS 2.1.0.2 patch package contains all content of the MRS 2.1.0.1 patch package.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## Impact of Patch Installation

- During the installation of the MRS 2.1.0.2 patch, MRS Manager will be restarted, and the components such as Hive, Spark, HDFS, YARN, MapReduce, Presto, HBase, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable, and services are not interrupted during the rolling restart.
- After installing the MRS 2.1.0.2 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).

- For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
- For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

 **NOTE**

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

## 7.8.10 MRS 2.1.0.1 Patch Description

### Basic Information

Table 7-45 Basic information

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>                    | MRS 2.1.0.1                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Date</b>                     | 2020-02-12                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Resolved Issues</b>                  | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b></p> <p>The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved.</p> <p><b>MRS big data components</b></p> <p>HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.</p> |
| <b>Compatibility with Other Patches</b> | None                                                                                                                                                                                                                                                                                                                                                                                                       |

### Impact of Patch Installation

During the installation of MRS 2.1.0.1 patches, MRS Manager and Hive are restarted. During the restart, the services are temporarily unavailable.

After MRS 2.1.0.1 patches are installed, log in to the Master1 node of the MRS cluster and delete the job directory in HDFS.

- For a cluster with Kerberos authentication disabled, run the following command to delete the job directory in HDFS:

```
hdfs dfs -rm -r /mrs/mrsjob/hive
```

- For a cluster with Kerberos authentication enabled, perform the following operations to delete the job directory in HDFS:
  - a. Run the following command and enter the password to perform authentication.  
**kinit hdfs**
  - b. Run the following command to delete the job directory in HDFS:  
**hdfs dfs -rm -r /mrs/mrsjob/hive**

 **NOTE**

This step is not required for a new MRS cluster because the directory does not exist in HDFS.

## 7.8.11 MRS 2.0.6.1 Patch Description

### Basic Information

**Table 7-46** Basic information

|                                         |                                                                                                                                                                                                                                                                        |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>                    | MRS 2.0.6.1                                                                                                                                                                                                                                                            |
| <b>Release Date</b>                     | 2020-07-06                                                                                                                                                                                                                                                             |
| <b>Resolved Issues</b>                  | <p><b>List of the resolved issues in MRS 2.1.0.1:</b></p> <p><b>MRS Manager</b><br/>Patch mechanism<br/>The monitoring metrics are empty occasionally.<br/>In DLF+Presto query, if a field contains a newline character, data and files are displayed incorrectly.</p> |
| <b>Compatibility with Other Patches</b> | None                                                                                                                                                                                                                                                                   |

### Impact of Patch Installation

During the installation of the MRS 2.0.6.1 patch, MRS Manager will be restarted, and the components such as Hive and services with dependency will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable, and services are not interrupted during the rolling restart.



## 7.8.12 MRS 2.0.1.3 Patch Description

### Basic Information

Table 7-47 Basic information

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>                    | MRS 2.0.1.3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Date</b>                     | 2019-12-25                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Resolved Issues</b>                  | <p><b>List of the resolved issues in MRS 2.0.1.3:</b></p> <p><b>MRS Manager</b></p> <p>The cluster scaling logic has been optimized and the TCP connection leak issues have been solved at the background of V1 job management APIs.</p> <p><b>MRS big data components</b></p> <p>The following issues have been solved for MRS Hive: HiveServer out of memory (OOM); slow <b>MergeFile</b> phase if a large number of small files exist; files not found in the <b>load partition</b> phase of <b>insert overwrite</b>, and file merging failure during <b>HIVE-22373:Container</b> reuse.</p> |
|                                         | <p><b>List of the resolved issues in MRS 2.0.1.2:</b></p> <p><b>MRS Manager</b></p> <p>The following issue has been solved: Scale-out fails occasionally due to timeout that occurs when ResourceManager executes <b>refreshNodes</b>.</p>                                                                                                                                                                                                                                                                                                                                                      |
|                                         | <p><b>List of the resolved issues in MRS 2.0.1.1:</b></p> <p><b>MRS Manager</b></p> <p>The following issue has been solved: OOM occurs on the executor of MRS Master nodes due to repeated node scale-in or scale-out.</p> <p><b>MRS big data components</b></p> <p>The following new function has been added: MRS Presto supports OBSFileSystem.</p> <p>The following issues have been solved for MRS Presto: The jstack is frequently printed and log files are too large to scroll.</p>                                                                                                      |
| <b>Compatibility with Other Patches</b> | The MRS 2.0.1.3 patch package contains all content of the MRS 2.0.1.2 and MRS 2.0.1.1 patch packages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## Impact of Patch Installation

During the installation of MRS 2.0.1.3 patches, MRS Manager and Presto are restarted. During the restart, the services are temporarily unavailable.

## 7.8.13 MRS 2.0.1.2 Patch Description

### Basic Information

Table 7-48 Basic information

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>                    | MRS 2.0.1.2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Date</b>                     | 2019-09-30                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Resolved Issues</b>                  | <p><b>List of the resolved issues in MRS 2.0.1.2:</b></p> <p><b>MRS Manager</b></p> <p>The following issue has been solved: Scale-out fails occasionally due to timeout that occurs when ResourceManager executes <b>refreshNodes</b>.</p> <p><b>List of the resolved issues in MRS 2.0.1.1:</b></p> <p><b>MRS Manager</b></p> <p>The following issue has been solved: OOM occurs on the executor of MRS Master nodes due to repeated node scale-in or scale-out.</p> <p><b>MRS big data components</b></p> <p>The following new function has been added: MRS Presto supports OBSFileSystem.</p> <p>The following issues have been solved for MRS Presto: The jstack is frequently printed and log files are too large to scroll.</p> |
| <b>Compatibility with Other Patches</b> | The MRS 2.0.1.2 patch package contains all content of the MRS 2.0.1.1 patch package.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Impact of Patch Installation

During the installation of MRS 2.0.1.2 patches, MRS Manager and Presto are restarted. During the restart, the services are temporarily unavailable.

## 7.8.14 MRS 2.0.1.1 Patch Description

### Basic Information

Table 7-49 Basic information

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>                    | MRS 2.0.1.1                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Date</b>                     | 2019-09-30                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Resolved Issues</b>                  | <p><b>MRS Manager</b></p> <p>The following issue has been solved: OOM occurs on the executor of MRS Master nodes due to repeated node scale-in or scale-out.</p> <p><b>MRS big data components</b></p> <p>The following new function has been added: MRS Presto supports OBSFileSystem.</p> <p>The following issues have been solved for MRS Presto: The jstack is frequently printed and log files are too large to scroll.</p> |
| <b>Compatibility with Other Patches</b> | None                                                                                                                                                                                                                                                                                                                                                                                                                             |

### Impact of Patch Installation

During the installation of MRS 2.0.1.1 patches, MRS Manager and Presto are restarted. During the restart, the services are temporarily unavailable.

## 7.8.15 MRS 1.9.3.3 Patch Description

### Basic Information

Table 7-50 Basic information

|                      |             |
|----------------------|-------------|
| <b>Patch Version</b> | MRS 1.9.3.3 |
| <b>Release Date</b>  | 2021-01-04  |

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Resolved Issues</b>                  | <p><b>List of resolved issues in MRS 1.9.3.3:</b></p> <p><b>MRS Manager</b><br/>Resolved the node isolation problem.</p> <p><b>MRS big data components</b><br/>Resolved the memory leak issue when Hive loads hooks.</p>                                                                                                                                                                                                                                                                                                                   |
|                                         | <p><b>List of resolved issues in MRS 1.9.3.2:</b></p> <p><b>MRS big data components</b><br/>When the insert overwrite operation is performed using Spark SQL and Beeline, old files cannot be trashed.</p>                                                                                                                                                                                                                                                                                                                                 |
|                                         | <p><b>List of resolved issues in MRS 1.9.3.1:</b></p> <p><b>MRS Manager</b><br/>Solved the problem that Task nodes fail to be removed from a custom cluster.</p> <p><b>MRS big data components</b><br/>Solved the problem that the version of the <b>adapter-hadoop-wrapper-file-system</b> package in the Hive and Spark paths is incorrect.</p> <p>Solved the problem that multiple namespaces saved on FusionInsight Manager of HBase do not take effect in the background.</p> <p>Added HDFSWrapper to support AbstractFileSystem.</p> |
| <b>Compatibility with Other Patches</b> | The MRS 1.9.3.3 patch package contains all patches released for MRS 1.9.3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Impact of Patch Installation

- During the installation of the MRS 1.9.3.3 patch, MRS Manager is restarted, and Hadoop, HDFS, Hive, Spark, and related dependent services are restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.3.3 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

 NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

## 7.8.16 MRS 1.9.3.1 Patch Description

### Basic Information

Table 7-51 Basic information

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>                    | MRS 1.9.3.1                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Date</b>                     | 2020-09-04                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Resolved Issues</b>                  | <p><b>MRS Manager</b><br/>Solved the problem that Task nodes fail to be removed from a custom cluster.</p> <p><b>MRS big data components</b><br/>Solved the problem that the version of the <b>adapter-hadoop-wrapper-file-system</b> package in the Hive and Spark paths is incorrect.</p> <p>Solved the problem that multiple namespaces saved on FusionInsight Manager of HBase do not take effect in the background.</p> <p>Added HDFSWrapper to support AbstractFileSystem.</p> |
| <b>Compatibility with Other Patches</b> | None                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

### Impact of Patch Installation

- During the installation of the MRS 1.9.3.1 patch, MRS Manager is restarted, and Hadoop, HDFS, Hive, Spark, and related dependent services are restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.3.1 patch, you need to download and install all clients again, including the original clients of the Master node and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

 NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

## 7.8.17 MRS 1.9.2.2 Patch Description

### Basic Information

Table 7-52 Basic information

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>   | MRS 1.9.2.2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Date</b>    | 2021-05-18                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Resolved Issues</b> | <p><b>MRS Manager</b><br/>Resolved the sudo privilege escalation vulnerability.<br/>Resolved the issue of queue information loss due to queue update during capacity expansion.</p> <p><b>MRS big data components</b><br/>Resolved the issue that the Hive on Spark task is suspended because the block ID is displayed as garbled characters.<br/>Self-developed APIs are added to Hive.<br/>Resolved the issue that the <b>map.xml</b> file cannot be read.<br/>Optimized the Hive Har feature.<br/>Resolved the issue that YARN is unavailable due to ZooKeeper dirty data.<br/>Upgraded the OBS packages.<br/>Upgraded the JDK version.<br/>Resolved the issue of ResourceManager memory leakage.<br/>Added the monitoring on the exception that occurs when the ECS getSecuritykey API is called.<br/>Optimized the temporary AK/SK process.<br/>Resolved the issue of ResourceManager memory leakage.<br/>Fixed the error reported when the Hive union statement is used to merge small files.<br/>Resolved the issue that the Hadoop task fails to be executed due to insufficient space.<br/>Resolved the issue that no data is generated after a Hive job is successfully executed.</p> |

|                                         |      |
|-----------------------------------------|------|
| <b>Compatibility with Other Patches</b> | None |
|-----------------------------------------|------|

## Impact of Patch Installation

- During the installation of the MRS 1.9.2.2 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Kafka, Ranger, Presto, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After the MRS 1.9.2.2 patch is installed, you need to restart the OMS service.

### NOTE

- Log in to the active and standby OMS nodes as user **root**, switch to user **omm**, and run the `sh ${BIGDATA_HOME}/om-0.0.1/sbin/restart-oms.sh` command to restart the OMS service.
- Both the active and standby OMS nodes need to be restarted.
- After installing the MRS 1.9.2.2 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

### NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

## 7.8.18 MRS 1.9.0.8, 1.9.0.9, and 1.9.0.10 Patch Description

### Basic Information

Table 7-53 Basic information

|                        |                                                                                                                                                                      |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Resolved Issues</b> | Patch number: <b>MRS 1.9.0.10</b><br>Release date: January 17, 2023<br><b>Resolved issues</b><br><b>MRS big data components</b><br>OBSA supports flow control retry. |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Patch number: <b>MRS 1.9.0.9</b><br/>Release date: August 10, 2022<br/><b>Resolved issues</b><br/><b>MRS big data components</b><br/>Superior scheduling algorithm optimization</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|  | <p>Patch number: <b>MRS 1.9.0.8</b><br/>Release date: February 20, 2021<br/><b>Resolved issues</b><br/><b>MRS big data components</b><br/>Added the monitoring on the exception that occurs when the ECS getSecuritykey API is called.<br/>Optimized the temporary AK/SK process.<br/>Resolved the issue of ResourceManager memory leakage.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | <p><b>List of resolved issues in MRS 1.9.0.7:</b><br/><b>MRS Manager</b><br/>Resolved the issue of queue information loss due to queue update during capacity expansion.<br/><b>MRS big data components</b><br/>Resolved the issue that the Hive on Spark task is suspended because the block ID is displayed as garbled characters.<br/>Solve the problem that the Hadoop task fails to be executed due to insufficient space.<br/>Self-developed APIs are added to Hive.<br/>Resolved the issue that the <b>map.xml</b> file cannot be read.<br/>Resolved the issue that YARN is unavailable due to ZooKeeper dirty data.<br/>Resolved the issue of ResourceManager memory leakage.<br/>Optimized the Hive Har feature.<br/>Upgraded the OBS packages.<br/>Upgraded the JDK version.</p> |



|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 1.9.0.6:</b></p> <p><b>MRS Manager</b><br/>MRS Manager supports scale-in of specified nodes in a yearly/monthly cluster.</p> <p><b>MRS big data components</b><br/>Resolved the issue of slow response when HiveSE delivers SQL statements.<br/>Supported the JobHistory query failure information interface.<br/>Resolved the issue that fine-grained permissions do not take effect.<br/>Fixed the exception that occurs when Hive on Spark reads data.<br/>Resolved the issue that data volume increases when the Hive on mrs task is executed twice.<br/>Resolved the issue that the performance of some strings is poor when vector-based vectorized query is enabled in Hive.</p> |
|--|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 1.9.0.5:</b></p> <p><b>MRS Manager</b></p> <p>Optimized the service restart process during configuration saving on MRS Manager.</p> <p>Rectified the periodic backup failure on MRS Manager.</p> <p><b>MRS big data components</b></p> <p>Private patch of Ranger</p> <p>Resolved the JVM Create GC thread failed issue of YARN.</p> <p>Added the HiveServer2 task stacking alarm.</p> <p>Added the alarm indicating that the GC time of HiveServer HiveMetastore exceeds 5s.</p> <p>Added the alarm indicating that HiveServer2 uncomment ZooKeeper.</p> <p>Added the alarm indicating that YARN tasks failed and the number of killed tasks exceeds 5 within 20 minutes.</p> <p>Corrected time zone of Spark JobHistory.</p> <p>Optimized MetaStore restart mechanism.</p> <p>Resolved the HIVE-22771 open-source issue.</p> <p>Resolved the Hive beeline log printing errors.</p> <p>Corrected the number of active nodes displayed on the YARN page.</p> <p>Resolved the slow response of RM page display, which is caused by large number of RM threads.</p> <p>Supported OBS monitoring.</p> <p>Upgraded the OBS packages.</p> <p>Resolved the issue that some data is not inserted when 10 data records are concurrently inserted into hive-jdbc.</p> <p>Resolved the issue that Hive occasionally reports a Kryo deserialization failure.</p> <p>Resolved the issue of Spark JobHistory memory leakage.</p> <p>Resolved the issue that the application list cannot be displayed occasionally in Spark JobHistory.</p> |
|  | <p><b>List of resolved issues in MRS 1.9.0.3:</b></p> <p><b>MRS Manager</b></p> <p>Upgraded Arm JDK on MRS Manager.</p> <p>Resolved the issue that the system disk is fully occupied by logs of the Core node on MRS Manager.</p> <p><b>MRS big data components</b></p> <p>Resolved the issue that the number of Ranger logs cannot be set, which may cause full disk occupation.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of resolved issues in MRS 1.9.0.2:</b></p> <p><b>MRS Manager</b><br/>Resolved the mutual trust lost between some Core nodes in the cluster.<br/>Resolved the issue that instances fail to be added after the patch is installed.<br/>Resolved the issue that the rolling restart timeout interval of HiveServer cannot be modified on MRS Manager.</p> <p><b>MRS big data components</b><br/>Upgraded the OBS packages.</p>                                                                                                            |
|                                                                 | <p><b>List of resolved issues in MRS 1.9.0.1:</b></p> <p><b>MRS Manager</b><br/>Resolved the issue that MRS Manager does not support rolling patch installation without restarting services.</p> <p><b>MRS big data components</b><br/>Resolved the issue that the OBS entrusted access frequency is not limited to 140 times within 5 minutes.<br/>Resolved the issue that Kafka does not support open-source access.<br/>Resolved the SPARK-27637 open-source issue.<br/>Optimized the Hive rolling restart.<br/>Upgraded the OBS packages.</p> |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 1.9.0.10 patch can resolve all problems that have been resolved by the MRS 1.9.0 patch.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## Impact of Patch Installation

- During the installation of the MRS 1.9.0.10 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Presto, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable, and services are not interrupted during the rolling restart.
- After installing the MRS 1.9.0.10 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

 NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) In the scenario where a temporary AK/SK is obtained by using an agency to access OBS, configure the **fs.obs.auth.node-cache-short-circuit.enable** parameter to determine whether to allow access to the ECS metadata API, thereby determining whether to trigger ECS flow control.

The MRS cluster can access OBS using a temporary AK/SK obtained by an agency. The temporary AK/SK is obtained using the ECS metadata API. The ECS metadata API has a flow control threshold of 140 times within 5 minutes for a single node. After the flow control is triggered, the node is added to the blacklist and cannot call the metadata API again within 30 minutes. To prevent flow control, MRS provides the node-level cross-process cache service meta to cache temporary AK/SK.

Application scenario: YARN jobs such as Spark and Hadoop that access OBS using a temporary AK/SK obtained by an agency. This parameter is configured in the **core-site.xml** file on the client.

The default value is **true**. The YARN application process in the MRS cluster obtains the temporary AK/SK from the node-level cache service meta. If meta is abnormal, obtain the temporary AK/SK from the ECS metadata API.

If you do not want to directly access the ECS metadata API when the meta is abnormal, set this parameter to **false** to prevent the node from being added to the blacklist due to flow control.

## 7.8.19 MRS 1.9.0.7 Patch Description

### Basic Information

Table 7-54 Basic information

|               |             |
|---------------|-------------|
| Patch Version | MRS 1.9.0.7 |
| Release Date  | 2021-01-15  |

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Resolved Issues</b> | <b>List of resolved issues in MRS 1.9.0.7:</b><br><b>MRS Manager</b><br>Solved the problem of queue information loss due to queue update during capacity expansion.<br><b>MRS big data components</b><br>Solved the problem that the Hive on Spark task is suspended because the block ID is displayed as garbled characters.<br>Solve the problem that the Hadoop task fails to be executed due to insufficient space.<br>Self-developed APIs are added to Hive.<br>Solved the problem that the <b>map.xml</b> file cannot be read.<br>Resolved the issue that YARN is unavailable due to ZooKeeper dirty data.<br>Resolved the issue of ResourceManager memory leakage.<br>The Hive Har feature is optimized.<br>Upgraded the OBS packages.<br>The JDK version is upgraded. |
|                        | <b>List of resolved issues in MRS 1.9.0.6:</b><br><b>MRS Manager</b><br>MRS Manager supports scale-in of specified nodes in a yearly/monthly cluster.<br><b>MRS big data components</b><br>Solved the problem of slow response when HiveSE delivers SQL statements.<br>Supported the jobhistory query failure information interface.<br>Solved the problem that fine-grained permissions do not take effect.<br>Fixed the exception that occurs when Hive on Spark reads data.<br>Solved the problem that data volume increases when the Hive on mrs task is executed twice.<br>Solved the problem that the performance of some strings is poor when vector-based vectorized query is enabled in Hive.                                                                        |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 1.9.0.5:</b></p> <p><b>MRS Manager</b></p> <p>Optimized the service restart process during configuration saving on MRS Manager.</p> <p>Rectified the periodic backup failure on MRS Manager.</p> <p><b>MRS big data components</b></p> <p>Private patch of Ranger</p> <p>Resolved the JVM Create GC thread failed issue of YARN.</p> <p>Added the HiveServer2 task stacking alarm.</p> <p>Added the alarm indicating that the GC time of HiveServer HiveMetastore exceeds 5s.</p> <p>Added the alarm indicating that HiveServer2 uncomment ZooKeeper.</p> <p>Added the alarm indicating that YARN tasks failed and the number of killed tasks exceeds 5 within 20 minutes.</p> <p>Corrected time zone of Spark JobHistory.</p> <p>Optimized MetaStore restart mechanism.</p> <p>Resolved the HIVE-22771 open-source issue.</p> <p>Resolved the Hive beeline log printing errors.</p> <p>Corrected the number of active nodes displayed on the YARN page.</p> <p>Resolved the slow response of RM page display, which is caused by large number of RM threads.</p> <p>Supported OBS monitoring.</p> <p>Upgraded the OBS packages.</p> <p>Resolved the issue that some data is not inserted when 10 data records are concurrently inserted into hive-jdbc.</p> <p>Resolved the issue that Hive occasionally reports a Kryo deserialization failure.</p> <p>Resolved the issue of Spark JobHistory memory leakage.</p> <p>Resolved the issue that the application list cannot be displayed occasionally in Spark JobHistory.</p> |
|  | <p><b>List of resolved issues in MRS 1.9.0.3:</b></p> <p><b>MRS Manager</b></p> <p>Upgraded Arm JDK on MRS Manager.</p> <p>Resolved the issue that the system disk is fully occupied by logs of the Core node on MRS Manager.</p> <p><b>MRS big data components</b></p> <p>Resolved the issue that the number of Ranger logs cannot be set, which may cause full disk occupation.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of resolved issues in MRS 1.9.0.2:</b></p> <p><b>MRS Manager</b><br/>Resolved the mutual trust lost between some Core nodes in the cluster.<br/>Resolved the issue that instances fail to be added after the patch is installed.<br/>Resolved the issue that the rolling restart timeout interval of HiveServer cannot be modified on MRS Manager.</p> <p><b>MRS big data components</b><br/>Upgraded the OBS packages.</p>                                                                                                            |
|                                                                 | <p><b>List of resolved issues in MRS 1.9.0.1:</b></p> <p><b>MRS Manager</b><br/>Resolved the issue that MRS Manager does not support rolling patch installation without restarting services.</p> <p><b>MRS big data components</b><br/>Resolved the issue that the OBS entrusted access frequency is not limited to 140 times within 5 minutes.<br/>Resolved the issue that Kafka does not support open-source access.<br/>Resolved the SPARK-27637 open-source issue.<br/>Optimized the Hive rolling restart.<br/>Upgraded the OBS packages.</p> |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 1.9.0.7 patch can resolve all problems that have been resolved by the MRS 1.9.0 patch.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## Impact of Patch Installation

- During the installation of the MRS 1.9.0.7 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Kafka, Ranger, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.0.7 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

 **NOTE**

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

## 7.8.20 MRS 1.9.0.6 Patch Description

### Basic Information

**Table 7-55** Basic information

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>   | MRS 1.9.0.6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Date</b>    | 2020-05-20                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Resolved Issues</b> | <p><b>List of resolved issues in MRS 1.9.0.6:</b></p> <p><b>MRS Manager</b><br/>MRS Manager supports scale-in of specified nodes in a yearly/monthly cluster.</p> <p><b>MRS big data components</b><br/>Solved the problem of slow response when HiveSE delivers SQL statements.<br/>Supported the jobhistory query failure information interface.<br/>Solved the problem that fine-grained permissions do not take effect.<br/>Fixed the exception that occurs when Hive on Spark reads data.<br/>Solved the problem that data volume increases when the Hive on mrs task is executed twice.<br/>Solved the problem that the performance of some strings is poor when vector-based vectorized query is enabled in Hive.</p> |



|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p><b>List of resolved issues in MRS 1.9.0.5:</b></p> <p><b>MRS Manager</b></p> <p>Optimized the service restart process during configuration saving on MRS Manager.</p> <p>Rectified the periodic backup failure on MRS Manager.</p> <p><b>MRS big data components</b></p> <p>Private patch of Ranger</p> <p>Resolved the JVM Create GC thread failed issue of YARN.</p> <p>Added the HiveServer2 task stacking alarm.</p> <p>Added the alarm indicating that the GC time of HiveServer HiveMetastore exceeds 5s.</p> <p>Added the alarm indicating that HiveServer2 uncomment ZooKeeper.</p> <p>Added the alarm indicating that YARN tasks failed and the number of killed tasks exceeds 5 within 20 minutes.</p> <p>Corrected time zone of Spark JobHistory.</p> <p>Optimized MetaStore restart mechanism.</p> <p>Resolved the HIVE-22771 open-source issue.</p> <p>Resolved the Hive beeline log printing errors.</p> <p>Corrected the number of active nodes displayed on the YARN page.</p> <p>Resolved the slow response of RM page display, which is caused by large number of RM threads.</p> <p>Supported OBS monitoring.</p> <p>Upgraded the OBS packages.</p> <p>Resolved the issue that some data is not inserted when 10 data records are concurrently inserted into hive-jdbc.</p> <p>Resolved the issue that Hive occasionally reports a Kryo deserialization failure.</p> <p>Resolved the issue of Spark JobHistory memory leakage.</p> <p>Resolved the issue that the application list cannot be displayed occasionally in Spark JobHistory.</p> |
|  | <p><b>List of resolved issues in MRS 1.9.0.3:</b></p> <p><b>MRS Manager</b></p> <p>Upgraded Arm JDK on MRS Manager.</p> <p>Resolved the issue that the system disk is fully occupied by logs of the Core node on MRS Manager.</p> <p><b>MRS big data components</b></p> <p>Resolved the issue that the number of Ranger logs cannot be set, which may cause full disk occupation.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of resolved issues in MRS 1.9.0.2:</b></p> <p><b>MRS Manager</b><br/>Resolved the mutual trust lost between some Core nodes in the cluster.<br/>Resolved the issue that instances fail to be added after the patch is installed.<br/>Resolved the issue that the rolling restart timeout interval of HiveServer cannot be modified on MRS Manager.</p> <p><b>MRS big data components</b><br/>Upgraded the OBS packages.</p>                                                                                                            |
|                                                                 | <p><b>List of resolved issues in MRS 1.9.0.1:</b></p> <p><b>MRS Manager</b><br/>Resolved the issue that MRS Manager does not support rolling patch installation without restarting services.</p> <p><b>MRS big data components</b><br/>Resolved the issue that the OBS entrusted access frequency is not limited to 140 times within 5 minutes.<br/>Resolved the issue that Kafka does not support open-source access.<br/>Resolved the SPARK-27637 open-source issue.<br/>Optimized the Hive rolling restart.<br/>Upgraded the OBS packages.</p> |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 1.9.0.6 patch can resolve all problems that have been resolved by the MRS 1.9.0 patch.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## Impact of Patch Installation

- During the installation of the MRS 1.9.0.6 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Kafka, Ranger, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.0.6 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

 **NOTE**

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

## 7.8.21 MRS 1.9.0.5 Patch Description

### Basic Information

**Table 7-56** Basic information

|                      |             |
|----------------------|-------------|
| <b>Patch Version</b> | MRS 1.9.0.5 |
| <b>Release Date</b>  | 2020-03-21  |

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Resolved Issues</b></p> | <p><b>List of resolved issues in MRS 1.9.0.5:</b></p> <p><b>MRS Manager</b></p> <p>Optimized the service restart process during configuration saving on MRS Manager.</p> <p>Rectified the periodic backup failure on MRS Manager.</p> <p><b>MRS big data components</b></p> <p>Private patch of Ranger</p> <p>Resolved the JVM Create GC thread failed issue of YARN.</p> <p>Added the HiveServer2 task stacking alarm.</p> <p>Added the alarm indicating that the GC time of HiveServer HiveMetastore exceeds 5s.</p> <p>Added the alarm indicating that HiveServer2 uncomment ZooKeeper.</p> <p>Added the alarm indicating that YARN tasks failed and the number of killed tasks exceeds 5 within 20 minutes.</p> <p>Corrected time zone of Spark JobHistory.</p> <p>Optimized MetaStore restart mechanism.</p> <p>Resolved the HIVE-22771 open-source issue.</p> <p>Resolved the Hive beeline log printing errors.</p> <p>Corrected the number of active nodes displayed on the YARN page.</p> <p>Resolved the slow response of RM page display, which is caused by large number of RM threads.</p> <p>Supported OBS monitoring.</p> <p>Upgraded the OBS packages.</p> <p>Resolved the issue that some data is not inserted when 10 data records are concurrently inserted into hive-jdbc.</p> <p>Resolved the issue that Hive occasionally reports a Kryo deserialization failure.</p> <p>Resolved the issue of Spark JobHistory memory leakage.</p> <p>Resolved the issue that the application list cannot be displayed occasionally in Spark JobHistory.</p> |
|                               | <p><b>List of resolved issues in MRS 1.9.0.3:</b></p> <p><b>MRS Manager</b></p> <p>Upgraded Arm JDK on MRS Manager.</p> <p>Resolved the issue that the system disk is fully occupied by logs of the Core node on MRS Manager.</p> <p><b>MRS big data components</b></p> <p>Resolved the issue that the number of Ranger logs cannot be set, which may cause full disk occupation.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

|                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                 | <p><b>List of resolved issues in MRS 1.9.0.2:</b></p> <p><b>MRS Manager</b><br/>Resolved the mutual trust lost between some Core nodes in the cluster.<br/>Resolved the issue that instances fail to be added after the patch is installed.<br/>Resolved the issue that the rolling restart timeout interval of HiveServer cannot be modified on MRS Manager.</p> <p><b>MRS big data components</b><br/>Upgraded the OBS packages.</p>                                                                                                            |
|                                                                 | <p><b>List of resolved issues in MRS 1.9.0.1:</b></p> <p><b>MRS Manager</b><br/>Resolved the issue that MRS Manager does not support rolling patch installation without restarting services.</p> <p><b>MRS big data components</b><br/>Resolved the issue that the OBS entrusted access frequency is not limited to 140 times within 5 minutes.<br/>Resolved the issue that Kafka does not support open-source access.<br/>Resolved the SPARK-27637 open-source issue.<br/>Optimized the Hive rolling restart.<br/>Upgraded the OBS packages.</p> |
| <p><b>Compati<br/>bility<br/>with<br/>Other<br/>Patches</b></p> | <p>The MRS 1.9.0.5 patch can resolve all problems that have been resolved by the MRS 1.9.0 patch.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## Impact of Patch Installation

- During the installation of the MRS 1.9.0.5 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Kafka, Ranger, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.0.5 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For how to fully update the original client of the active and standby master nodes, see [Updating Client Configurations \(Version 2.x or Earlier\)](#).
  - For details about how to fully install the clients you set up, see [Installing a Client \(MRS 2.x or Earlier\)](#).

 NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

## 7.8.22 MRS 1.8.10.1 Patch Description

### Basic Information

Table 7-57 Basic information

|                                         |                                                                                                                |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Patch Version</b>                    | MRS 1.8.10.1                                                                                                   |
| <b>Release Date</b>                     | 2020-01-07                                                                                                     |
| <b>Resolved Issues</b>                  | <b>MRS big data components</b><br>The health check and rolling restart logic of MRS Kafka have been optimized. |
| <b>Compatibility with Other Patches</b> | None                                                                                                           |

### Impact of Patch Installation

During the installation of MRS 1.8.10.1 patches, MRS Manager and Kafka are restarted. During the restart, the services are temporarily unavailable.

## 7.9 Viewing Logs of an MRS Cluster

### 7.9.1 Overview of MRS Cluster Logs

#### Log Description

MRS cluster logs are stored in the `/var/log/Bigdata` directory. The following table lists the log types.

**Table 7-58** Log types

| Type             | Description                                                                                                                                                                                       |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Installation log | Installation logs record information about FusionInsight Manager, cluster, and service installation to help users locate installation errors.                                                     |
| Run logs         | Run logs record the running track information, debugging information, status changes, potential problems, and error information generated during the running of services.                         |
| Audit logs       | Audit logs record information about users' activities and operation instructions, which can be used to locate fault causes in security events and determine who are responsible for these faults. |

The following table lists the MRS log directories.

**Table 7-59** Log directories

| File Directory              | Log Content                                                                         |
|-----------------------------|-------------------------------------------------------------------------------------|
| /var/log/Bigdata/audit      | Component audit log.                                                                |
| /var/log/Bigdata/controller | Log collecting script log.<br>Controller process log.<br>Controller monitoring log. |
| /var/log/Bigdata/dbservice  | DBService log.                                                                      |
| /var/log/Bigdata/flume      | Flume log.                                                                          |
| /var/log/Bigdata/hbase      | HBase log.                                                                          |
| /var/log/Bigdata/hdfs       | HDFS log.                                                                           |
| /var/log/Bigdata/hive       | Hive log.                                                                           |
| /var/log/Bigdata/hetuengine | HetuEngine logs.                                                                    |
| /var/log/Bigdata/httpd      | HTTPD log.                                                                          |
| /var/log/Bigdata/hue        | Hue log.                                                                            |
| /var/log/Bigdata/kerberos   | Kerberos log.                                                                       |
| /var/log/Bigdata/ldapclient | LDAP client log.                                                                    |
| /var/log/Bigdata/ldapserver | LDAP server log.                                                                    |
| /var/log/Bigdata/loader     | Loader log.                                                                         |
| /var/log/Bigdata/logman     | logman script log management log.                                                   |
| /var/log/Bigdata/mapreduce  | MapReduce log.                                                                      |

| File Directory                  | Log Content                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /var/log/Bigdata/nodeagent      | NodeAgent log.                                                                                                                                                                                                                                                                                                                   |
| /var/log/Bigdata/okerberos      | OMS Kerberos log.                                                                                                                                                                                                                                                                                                                |
| /var/log/Bigdata/oldapserver    | OMS LDAP log.                                                                                                                                                                                                                                                                                                                    |
| /var/log/Bigdata/metric_agent   | MetricAgent run log.                                                                                                                                                                                                                                                                                                             |
| /var/log/Bigdata/omm            | <b>oms</b> : complex event processing log, alarm service log, HA log, authentication and authorization management log, and monitoring service run log of the omm server.<br><b>oma</b> : installation log and run log of the omm agent.<br><b>core</b> : dump log generated when the omm agent and the HA process are suspended. |
| /var/log/Bigdata/spark2x        | Spark2x log.                                                                                                                                                                                                                                                                                                                     |
| /var/log/Bigdata/sudo           | Log generated when the <b>sudo</b> command is executed by user <b>omm</b> .                                                                                                                                                                                                                                                      |
| /var/log/Bigdata/timestamp      | Time synchronization management log.                                                                                                                                                                                                                                                                                             |
| /var/log/Bigdata/tomcat         | Tomcat log.                                                                                                                                                                                                                                                                                                                      |
| /var/log/Bigdata/watchdog       | Watchdog log.                                                                                                                                                                                                                                                                                                                    |
| /var/log/Bigdata/yarn           | YARN log.                                                                                                                                                                                                                                                                                                                        |
| /var/log/Bigdata/zookeeper      | ZooKeeper log.                                                                                                                                                                                                                                                                                                                   |
| /var/log/Bigdata/oozie          | Oozie log.                                                                                                                                                                                                                                                                                                                       |
| /var/log/Bigdata/kafka          | Kafka log.                                                                                                                                                                                                                                                                                                                       |
| /var/log/Bigdata/storm          | Storm log.                                                                                                                                                                                                                                                                                                                       |
| /var/log/Bigdata/iotdb          | IoTDB log.                                                                                                                                                                                                                                                                                                                       |
| /var/log/Bigdata/cdl            | CDL log.                                                                                                                                                                                                                                                                                                                         |
| /var/log/Bigdata/upgrade        | OMS upgrade log.                                                                                                                                                                                                                                                                                                                 |
| /var/log/Bigdata/update-service | Upgrade service log.                                                                                                                                                                                                                                                                                                             |
| /var/log/Bigdata/patch          | Patch log.                                                                                                                                                                                                                                                                                                                       |

## Run logs

[Table 7-60](#) describes the running information recorded in run logs.



**Table 7-60** Running information

| Run Log                         | Description                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Installation preparation log    | Records information about preparations for the installation, such as the detection, configuration, and feedback operation information.                                    |
| Process startup log             | Records information about the commands executed during the process startup.                                                                                               |
| Process startup exception log   | Records information about exceptions during process startup, such as dependent service errors and insufficient resources.                                                 |
| Process run log                 | Records information about the process running track information and debugging information, such as function entries and exits as well as cross-module interface messages. |
| Process running exception log   | Records errors that cause process running errors, for example, the empty input objects or encoding or decoding failure.                                                   |
| Process running environment log | Records information about the process running environment, such as resource status and environment variables.                                                             |
| Script logs                     | Records information about the script execution process.                                                                                                                   |
| Resource reclamation log        | Records information about the resource reclaiming process.                                                                                                                |
| Uninstallation clearing logs    | Records information about operations performed during service uninstallation, such as directory deletion and execution time                                               |

## Audit logs

Audit information recorded in audit logs includes FusionInsight Manager audit information and component audit information.

**Table 7-61** Audit information of FusionInsight Manager

| Operation Type  | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| User management | Creating a user<br>Modifying a user<br>Deleting a user<br>Creating a user group<br>Modifying a user group<br>Deleting a user group<br>Adding a role<br>Modifying a role<br>Deleting a role<br>Changing a password policy<br>Changing a password<br>Resetting a password<br>User login<br>User logout<br>Unlocking the screen<br>Downloading the authentication credential<br>Unauthorized operation<br>Unlocking a user account<br>Locking a user account<br>Locking the screen<br>Exporting user information<br>Exporting a user group<br>Exporting a role |

| Operation Type | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster        | <ul style="list-style-type: none"> <li>Starting a cluster</li> <li>Stopping a cluster</li> <li>Restarting a cluster</li> <li>Performing a rolling restart of a cluster</li> <li>Restarting all expired instances</li> <li>Saving configurations</li> <li>Synchronizing cluster configurations</li> <li>Customizing cluster monitoring indicators</li> <li>Configuring monitoring dumping</li> <li>Saving monitoring thresholds</li> <li>Downloading a client configuration file</li> <li>Configuring the northbound Syslog interface</li> <li>Configuring the northbound SNMP API</li> <li>Clearing alarm through SNMP</li> <li>Adding a trap target through SNMP</li> <li>Deleting a trap target through SNMP</li> <li>Checking alarms through SNMP</li> <li>Synchronizing alarms through SNMP</li> <li>Creating a threshold template</li> <li>Deleting a threshold template</li> <li>Applying a threshold template</li> <li>Saving cluster monitoring configuration data</li> <li>Exporting configuration data</li> <li>Importing cluster configuration data</li> <li>Exporting an installation template</li> <li>Modifying a threshold template</li> <li>Canceling the application of a threshold template</li> <li>Masking alarms</li> <li>Sending an alarm</li> <li>Changing the OMS database password</li> <li>Resetting the component database password</li> <li>Restarting OMM and Controller</li> <li>Starting the health check of a cluster</li> <li>Importing a certificate file</li> <li>Configuring SSO information</li> <li>Deleting historical health check reports</li> <li>Modifying cluster attributes</li> </ul> |

| Operation Type | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                | Running maintenance commands in synchronous mode<br>Running maintenance commands in asynchronous mode<br>Customizing report monitoring indicators<br>Exporting report monitoring data<br>Running a command using SNMP in asynchronous mode<br>Restarting the Web service<br>Customizing monitoring indicators for static resource pools<br>Exporting monitoring data of a static resource pool<br>Customizing dashboard monitoring metrics<br>Stopping a task<br>Restoring configurations<br>Modifying domain and mutual trust configurations<br>Modifying system parameters<br>Making a cluster enter the maintenance mode<br>Making a cluster exit the maintenance mode<br>Making OMS enter the maintenance mode<br>Making OMS exit the maintenance mode<br>Exiting the maintenance mode in batches<br>Modifying OMS configurations<br>Enabling threshold alarms<br>Synchronizing all cluster configurations |

| Operation Type | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Service        | <ul style="list-style-type: none"> <li>Starting a service</li> <li>Stopping a service</li> <li>Synchronizing service configurations</li> <li>Refreshing a service queue</li> <li>Customizing service monitoring indicators</li> <li>Restarting a service</li> <li>Performing a rolling restart of a service</li> <li>Exporting service monitoring data</li> <li>Importing service configuration data</li> <li>Starting the health check of a service</li> <li>Configuring the service</li> <li>Uploading a configuration file</li> <li>Downloading a configuration file</li> <li>Synchronizing instance configurations</li> <li>Commissioning an instance</li> <li>Decommissioning an instance</li> <li>Starting an instance</li> <li>Stopping an instance</li> <li>Customizing instance monitoring indicators</li> <li>Restarting an instance</li> <li>Performing a rolling restart of an instance</li> <li>Exporting instance monitoring data</li> <li>Importing instance configuration data</li> <li>Creating an instance group</li> <li>Modifying an instance group</li> <li>Deleting an instance group</li> <li>Moving an instance to another instance group</li> <li>Making a service enter the maintenance mode</li> <li>Making a service exit the maintenance mode</li> <li>Changing the display name of a service</li> <li>Modifying service association</li> <li>Downloading monitoring data</li> <li>Masking alarms</li> <li>Unmasking alarms</li> <li>Exporting report data of a service</li> <li>Adding custom parameters for a report</li> <li>Modifying custom parameters of a report</li> <li>Deleting custom parameters of a report</li> </ul> |

| Operation Type | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                | Switching over controller nodes<br>Adding a mount table<br>Modifying a mount table                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Host           | Setting a node rack<br>Starting all roles<br>Stopping all roles<br>Isolating a host<br>Canceling host isolation<br>Customizing host monitoring indicators<br>Exporting host monitoring data<br>Making a host enter the maintenance mode<br>Making a host exit the maintenance mode<br>Exporting basic host information<br>Exporting host distribution report data<br>Exporting host trend report data<br>Exporting host cluster report data<br>Exporting service report data<br>Customizing host cluster monitoring metrics<br>Customizing host trend monitoring metrics |
| Alarm          | Exporting alarms<br>Clearing alarms<br>Exporting events<br>Clearing alarms in batches                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Log collection | Collecting log files<br>Downloading log files<br>Collecting service stack information<br>Collecting instance stack information<br>Preparing service stack information<br>Preparing instance stack information<br>Clearing service stack information<br>Clearing instance stack information                                                                                                                                                                                                                                                                               |
| Audit logs     | Modifying audit dump configurations<br>Exporting audit logs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

| Operation Type         | Operation                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup and restoration | Creating a backup task<br>Executing a backup task<br>Executing backup tasks in batches<br>Stopping a backup task<br>Deleting a backup task<br>Modifying a backup task<br>Locking a backup task<br>Unlocking a backup task<br>Creating a restoration task<br>Executing a backup restoration task<br>Stopping a restoration task<br>Retrying a restoration task<br>Deleting a restoration task |

| Operation Type          | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Multi-tenant management | Saving the static configuration<br>Creating a tenant<br>Deleting a tenant<br>Associating a service with a tenant<br>Deleting a service from a tenant<br>Configuring resources<br>Creating resources<br>Deleting resources<br>Adding a resource pool<br>Modifying a resource pool<br>Deleting a resource pool<br>Restoring tenant data<br>Modifying a tenant's global configurations<br>Modifying queue configurations for a capacity scheduler<br>Modifying queue configurations for a super scheduler<br>Modifying resource distribution for a capacity scheduler<br>Clearing resource distribution for a capacity scheduler<br>Modifying resource distribution for a super scheduler<br>Clearing resource distribution for a super scheduler<br>Adding a resource catalog<br>Modifying a resource catalog<br>Deleting a resource catalog<br>Customizing tenant monitoring metrics |



| Operation Type | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Health check   | Starting health check for a cluster<br>Starting health check for a service<br>Starting health check for a host<br>Starting health check for OMS<br>Starting health check for a system<br>Updating health check configurations<br>Exporting health check reports<br>Exporting cluster health check results<br>Exporting service health check results<br>Exporting host health check results<br>Deleting historical health check reports<br>Exporting historical health check reports<br>Downloading health check reports |

**Table 7-62** Component audit information

| Audit Log            | Operation Type         | Operation                                                                                      |
|----------------------|------------------------|------------------------------------------------------------------------------------------------|
| CDL audit log        | Service operation      | Creating a link<br>Deleting a link<br>Creating a job<br>Starting a job<br>Deleting a job       |
| IoTDB audit log      | Maintenance management | Granting permissions<br>Revoking permissions<br>Recording authentication and login information |
|                      | Service operation      | Deleting a time series, partition, function, or index<br>Modifying a time series               |
| ClickHouse audit log | Maintenance management | Granting permissions<br>Revoking permissions<br>Recording authentication and login information |

| Audit Log           | Operation Type                             | Operation                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     | Service operation                          | Creating a database or table<br>Inserting, deleting, querying, and migrating data                                                                                                                                                                                                                                                                                                                    |
| DBService audit log | Maintenance management                     | Performing backup restoration operations                                                                                                                                                                                                                                                                                                                                                             |
| HBase audit log     | Data definition language (DDL) statement   | Creating a table<br>Deleting a table<br>Modifying a table<br>Adding a column family<br>Modifying a column family<br>Deleting a column family<br>Enabling a table<br>Disabling a table<br>Modify the user information<br>Changing a password<br>User login                                                                                                                                            |
|                     | Data manipulation language (DML) statement | Putting data (to the <b>hbase:meta</b> , <b>_ctmeta</b> , and <b>hbase:acl</b> tables)<br>Deleting data (from the <b>hbase:meta</b> , <b>_ctmeta</b> , and <b>hbase:acl</b> tables)<br>Checking and putting data (to the <b>hbase:meta</b> , <b>_ctmeta</b> , and <b>hbase:acl</b> tables)<br>Checking and deleting data (from the <b>hbase:meta</b> , <b>_ctmeta</b> , and <b>hbase:acl</b> tables) |
|                     | Permission control                         | Assigning permissions to a user<br>Canceling permission assigning                                                                                                                                                                                                                                                                                                                                    |

| Audit Log       | Operation Type         | Operation                                                                                                                                                                                                                                                                |
|-----------------|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HDFS audit log  | Permissions management | Managing permissions on files or folders<br>Managing permissions on owner information files or folders                                                                                                                                                                   |
|                 | File operation         | Creating a folder<br>Creating a file<br>Opening a file<br>Appending file content<br>Changing a file name<br>Deleting a file or folder<br>Setting time property of a file<br>Setting the number of file copies<br>Merging files<br>Checking the file system<br>File links |
| Hive audit logs | Metadata operation     | Defining metadata, such as creating databases and tables<br>Deleting metadata, such as deleting databases and tables<br>Modifying metadata, such as adding columns and renaming tables<br>Importing and exporting metadata                                               |
|                 | Data maintenance       | Loading data to a table<br>Inserting data into a table                                                                                                                                                                                                                   |
|                 | Permissions management | Creating or deleting roles<br>Granting/Reclaiming roles<br>Granting/Reclaiming permissions                                                                                                                                                                               |
| Hue audit log   | Service startup        | Starting Hue                                                                                                                                                                                                                                                             |
|                 | User operation         | User login<br>User logout                                                                                                                                                                                                                                                |

| Audit Log            | Operation Type                   | Operation                                                                                                                                                                                                                   |
|----------------------|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                      | Task operation                   | Creating a job<br>Modifying a job<br>Deleting a job<br>Submitting a task<br>Saving a task<br>Updating the status of a task                                                                                                  |
| KrbServer audit log  | Maintenance management           | Changing the password of a Kerberos account<br>Adding a Kerberos account<br>Deleting a Kerberos account<br>Authenticating a user                                                                                            |
| LdapServer audit log | Maintenance management           | Adding an operating system user<br>Adding a user group<br>Adding a user to user group<br>Deleting a user<br>Deleting a group                                                                                                |
| Loader audit log     | Security management              | User login                                                                                                                                                                                                                  |
|                      | Metadata management              | Querying connector information<br>Querying a framework<br>Querying step information                                                                                                                                         |
|                      | Managing data source connections | Querying a data source connection<br>Adding a data source connection<br>Updating a data source connection<br>Deleting a data source connection<br>Activating a data source connection<br>Disabling a data source connection |

| Audit Log           | Operation Type      | Operation                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     | Job management      | Querying a job<br>Creating a Job<br>Updating a Job<br>Deleting a job<br>Activating a job<br>Disabling a job<br>Querying all execution records of a job<br>Querying the latest execution record of a job<br>Submitting a job<br>Stopping a job                                                                                                                                             |
| MapReduce audit log | Application running | Starting a Container request<br>Stopping a Container request<br>After Container request is completed, the status of the request is displayed as succeeded.<br>After Container request is completed, the status of the request is displayed as failed.<br>After Container request is completed, the status of the request is displayed as suspended.<br>Submitting a task<br>Ending a task |
| Oozie audit log     | Task management     | Submitting a task<br>Starting a task<br>Killing a task<br>Suspending a task<br>Resuming a task<br>Re-running a task                                                                                                                                                                                                                                                                       |

| Audit Log           | Operation Type         | Operation                                                                                                                                                                                                                  |
|---------------------|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Spark audit log     | Metadata operation     | Defining metadata, such as creating databases and tables<br>Deleting metadata, such as deleting databases and tables<br>Modifying metadata, such as adding columns and renaming tables<br>Importing and exporting metadata |
|                     | Data maintenance       | Loading data to a table<br>Inserting data into a table                                                                                                                                                                     |
| Storm audit log     | Nimbus                 | Submitting a topology<br>Stopping a topology<br>Reallocating a topology<br>Deactivating a topology<br>Activating a topology                                                                                                |
|                     | UI                     | Stopping a topology<br>Reallocating a topology<br>Deactivating a topology<br>Activating a topology                                                                                                                         |
| YARN audit log      | Job submission         | Submitting a job to a queue                                                                                                                                                                                                |
| ZooKeeper audit log | Permissions management | Setting the access permission to Znode                                                                                                                                                                                     |
|                     | Znode operation        | Creating a Znode<br>Deleting a Znode<br>Configuring Znode data                                                                                                                                                             |

| Audit Log            | Operation Type | Operation                                                                                                                                                                                                                                                                                                       |
|----------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HetuEngine audit log | Job management | Adding an external data source<br>Deleting an external data source<br>Modifying an external data source<br>Creating a compute instance<br>Starting a compute instance<br>Stopping a compute instance<br>Deleting a compute instance<br>Querying a compute instance<br>Modifying compute instance configurations |

MRS audit logs are stored in the database. You can view and export audit logs on the **Audit** page.

The following table lists the directories to store component audit logs. Audit log files of some components are stored in **/var/log/Bigdata/audit**, such as HDFS, HBase, MapReduce, Hive, Hue, YARN, Storm, and ZooKeeper. The component audit logs are automatically compressed and backed up to **/var/log/Bigdata/audit/bk** at 03:00 every day. A maximum of latest 90 compressed backup files are retained, and the backup time cannot be changed.

Audit log files of other components are stored in the component log directory.

**Table 7-63** Directory for storing component audit logs

| Component | Audit Log Directory                                  |
|-----------|------------------------------------------------------|
| DBService | /var/log/Bigdata/audit/dbservice/dbservice_audit.log |

| Component  | Audit Log Directory                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HBase      | <p>/var/log/Bigdata/audit/hbase/hm/hbase-audit-hmaster.log</p> <p>/var/log/Bigdata/audit/hbase/hm/hbase-ranger-audit-hmaster.log</p> <p>/var/log/Bigdata/audit/hbase/rs/hbase-audit-regionserver.log</p> <p>/var/log/Bigdata/audit/hbase/rs/hbase-ranger-audit-regionserver.log</p> <p>/var/log/Bigdata/audit/hbase/rt/hbase-audit-restserver.log</p> <p>/var/log/Bigdata/audit/hbase/ts/hbase-audit-thriftserver.log</p>                                                                           |
| HDFS       | <p>/var/log/Bigdata/audit/hdfs/nn/hdfs-audit-namenode.log</p> <p>/var/log/Bigdata/audit/hdfs/nn/ranger-plugin-audit.log</p> <p>/var/log/Bigdata/audit/hdfs/dn/hdfs-audit-datanode.log</p> <p>/var/log/Bigdata/audit/hdfs/jn/hdfs-audit-journalnode.log</p> <p>/var/log/Bigdata/audit/hdfs/zkfc/hdfs-audit-zkfc.log</p> <p>/var/log/Bigdata/audit/hdfs/httpfs/hdfs-audit-httpfs.log</p> <p>/var/log/Bigdata/audit/hdfs/router/hdfs-audit-router.log</p>                                              |
| HetuEngine | <p>/var/log/Bigdata/audit/hetuengine/hsbroker/hsbroker-audit.log.0</p> <p>/var/log/Bigdata/audit/hetuengine/hsconsole/hsconsole-audit.log.0</p> <p>/var/log/Bigdata/audit/hetuengine/hsfabric/hsfabric-audit.log.0</p> <p>hdfs://hacluster/hetuserverhistory/<i>Tenant name</i>/coordinator/application_ID/container_ID/yyyyMMdd/hetuserver-engine-audit.log</p> <p>hdfs://hacluster/hetuserverhistory/<i>Tenant name</i>/coordinator or worker/application_ID/container_ID/yyyyMMdd/server.log</p> |



| Component | Audit Log Directory                                                                                                                                                                                                                                                                            |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hive      | <code>/var/log/Bigdata/audit/hive/hiveserver/hive-audit.log</code><br><code>/var/log/Bigdata/audit/hive/hiveserver/hive-rangeraudit.log</code><br><code>/var/log/Bigdata/audit/hive/metastore/metastore-audit.log</code><br><code>/var/log/Bigdata/audit/hive/webhcat/webhcat-audit.log</code> |
| Hue       | <code>/var/log/Bigdata/audit/hue/hue-audits.log</code>                                                                                                                                                                                                                                         |
| Kafka     | <code>/var/log/Bigdata/audit/kafka/audit.log</code>                                                                                                                                                                                                                                            |
| Loader    | <code>/var/log/Bigdata/loader/audit/default.audit</code>                                                                                                                                                                                                                                       |
| CDL       | <code>/var/log/Bigdata/audit/cdl/service/cdl-audit.log</code>                                                                                                                                                                                                                                  |
| MapReduce | <code>/var/log/Bigdata/audit/mapreduce/jobhistory/mapred-audit-jobhistory.log</code>                                                                                                                                                                                                           |
| Oozie     | <code>/var/log/Bigdata/audit/oozie/oozie-audit.log</code>                                                                                                                                                                                                                                      |
| Spark2x   | <code>/var/log/Bigdata/audit/spark2x/jdbcserver/jdbcserver-audit.log</code><br><code>/var/log/Bigdata/audit/spark2x/jdbcserver/ranger-audit.log</code><br><code>/var/log/Bigdata/audit/spark2x/jobhistory/jobhistory-audit.log</code>                                                          |
| Storm     | <code>/var/log/Bigdata/audit/storm/logviewer/audit.log</code><br><code>/var/log/Bigdata/audit/storm/nimbus/audit.log</code><br><code>/var/log/Bigdata/audit/storm/supervisor/audit.log</code><br><code>/var/log/Bigdata/audit/storm/ui/audit.log</code>                                        |
| Yarn      | <code>/var/log/Bigdata/audit/yarn/rm/yarn-audit-resource-manager.log</code><br><code>/var/log/Bigdata/audit/yarn/rm/ranger-plugin-audit.log</code><br><code>/var/log/Bigdata/audit/yarn/nm/yarn-audit-nodemanager.log</code>                                                                   |
| ZooKeeper | <code>/var/log/Bigdata/audit/zookeeper/quorumpeer/zk-audit-quorumpeer.log</code>                                                                                                                                                                                                               |
| IoTDB     | <code>/var/log/Bigdata/audit/iotdb/iotdbserver/log_audit.log</code>                                                                                                                                                                                                                            |

## Manager Logs

**Log path:** Manager log files are stored in `/var/log/Bigdata/Manager` by default.

- ControllerService: **/var/log/Bigdata/controller/** (OMS installation and run logs)
- Httpd: **/var/log/Bigdata/httpd** (httpd installation and run logs)
- logman: **/var/log/Bigdata/logman** (log packaging tool logs)
- NodeAgent: **/var/log/Bigdata/nodeagent** (NodeAgent installation and run logs)
- okerberos: **/var/log/Bigdata/okerberos** (okerberos installation and run logs)
- oldapserver: **/var/log/Bigdata/oldapserver** (oldapserver installation and run logs)
- MetricAgent: **/var/log/Bigdata/metric\_agent** (MetricAgent run logs)
- omm: **/var/log/Bigdata/omm** (omm installation and run logs)
- timestamp: **/var/log/Bigdata/timestamp** (NodeAgent startup time logs)
- tomcat: **/var/log/Bigdata/tomcat** (Web process logs)
- Watchdog: **/var/log/Bigdata/watchdog** (watchdog logs)
- Upgrade: **/var/log/Bigdata/upgrade** (OMS upgrade logs)
- UpdateService: **/var/log/Bigdata/update-service** (upgrade service logs)
- Sudo: **/var/log/Bigdata/sudo** (sudo script execution logs)
- OS: **/var/log/message file** (OS system logs)
- OS Performance: **/var/log/osperf** (OS performance statistics logs)
- OS Statistics: **/var/log/osinfo/statistics** (OS parameter configuration logs)

**Log archiving rules:**

Manager logs are automatically compressed and archived. When a log file exceeds 10 MB in size, the file is automatically compressed by default. The file is named as follows after compression: *<Original log name>-<yyyy-mm-dd\_hh-mm-ss>.[ID].log.zip*. A maximum of 20 latest compressed files are reserved.

**Table 7-64** Manager logs

| Type               | Log File Name         | Description                                                                                                                   |
|--------------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Controller run log | controller.log        | Log file that records component installation, upgrade, configuration, monitoring, alarm reporting, and routine O&M operations |
|                    | controller_client.log | Run log file of REST APIs                                                                                                     |
|                    | acs.log               | ACS run log file                                                                                                              |
|                    | acs_spnego.log        | spnego user log in ACS                                                                                                        |
|                    | aos.log               | AOS run log                                                                                                                   |
|                    | plugin.log            | AOS plug-in log                                                                                                               |

| Type | Log File Name                                                                                                                 | Description                                                                                     |
|------|-------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
|      | backupplugin.log                                                                                                              | Log file that records backup and restoration operations                                         |
|      | controller_config.log                                                                                                         | Configuration run log                                                                           |
|      | controller_nodesetup.log                                                                                                      | Controller loading task log                                                                     |
|      | controller_root.log                                                                                                           | System log of the Controller process                                                            |
|      | controller_trace.log                                                                                                          | Log that records the remote procedure call (RPC) communication between Controller and NodeAgent |
|      | controller_monitor.log                                                                                                        | Monitoring log                                                                                  |
|      | controller_fsm.log                                                                                                            | State machine log                                                                               |
|      | controller_alarm.log                                                                                                          | Controller alarm log                                                                            |
|      | controller_backup.log                                                                                                         | Controller backup and restoration log                                                           |
|      | install.log,<br>restore_package.log,<br>installPack.log,<br>distributeAdapterFiles.log,<br>and<br>install_os_optimization.log | OMS installation log                                                                            |
|      | oms_ctl.log                                                                                                                   | OMS startup and stop log                                                                        |
|      | preInstall_client.log                                                                                                         | Preprocessing log file before client installation                                               |
|      | installntp.log                                                                                                                | NTP installation log                                                                            |
|      | modify_manager_param.log                                                                                                      | Manager parameter modification log                                                              |
|      | backup.log                                                                                                                    | OMS backup script run log                                                                       |
|      | supressionAlarm.log                                                                                                           | Alarm script run log                                                                            |
|      | om.log                                                                                                                        | OM certificate generation log                                                                   |

| Type      | Log File Name                            | Description                                                                                                    |
|-----------|------------------------------------------|----------------------------------------------------------------------------------------------------------------|
|           | backupplugin_ctl.log                     | Startup log of the backup and restoration plug-in process                                                      |
|           | getLogs.log                              | Run log of the collection log script                                                                           |
|           | backupAuditLogs.log                      | Run log of the audit log backup script                                                                         |
|           | certStatus.log                           | Log that records regular certificate checks                                                                    |
|           | distribute.log                           | Certificate distribution log                                                                                   |
|           | ficertgenenerate.log                     | Certificate replacement logs, including logs of level-2 certificates, CAS certificates, and httpd certificates |
|           | genPwFile.log                            | Log that records the generation of certificate password files                                                  |
|           | modifyproxyconf.log                      | Log that records the modification of the HTTPD proxy configuration                                             |
|           | importTar.log                            | Log that records the process of importing certificates into the trust library                                  |
| Httpd     | install.log                              | Httpd installation log                                                                                         |
|           | access_log, error_log                    | Httpd run log                                                                                                  |
| logman    | logman.log                               | Log packaging tool log                                                                                         |
| NodeAgent | install.log, install_os_optimization.log | NodeAgent installation log                                                                                     |
|           | installntp.log                           | NTP installation log                                                                                           |
|           | start_ntp.log                            | NTP startup log                                                                                                |
|           | ntpChecker.log                           | NTP check log                                                                                                  |
|           | ntpMonitor.log                           | NTP monitoring log                                                                                             |

| Type      | Log File Name                            | Description                                                         |
|-----------|------------------------------------------|---------------------------------------------------------------------|
|           | heartbeat_trace.log                      | Log that records heartbeats between NodeAgent and Controller        |
|           | alarm.log                                | Alarm log                                                           |
|           | monitor.log                              | Monitoring log                                                      |
|           | nodeagent_ctl.log, start-agent.log       | NodeAgent startup log                                               |
|           | agent.log                                | NodeAgent run log                                                   |
|           | cert.log                                 | Certificate log                                                     |
|           | agentplugin.log                          | Agent plug-in running status monitoring log                         |
|           | omapplugin.log                           | OMA plug-in run log                                                 |
|           | diskhealth.log                           | Disk health check log                                               |
|           | supressionAlarm.log                      | Alarm script run log                                                |
|           | updateHostFile.log                       | Host list update log                                                |
|           | collectLog.log                           | Run log of the node log collection script                           |
|           | host_metric_collect.log                  | Host index collection run log                                       |
|           | checkfileconfig.log                      | Run log file of file permission check                               |
|           | entropycheck.log                         | Entropy check run log                                               |
|           | timer.log                                | Log of periodic node scheduling                                     |
|           | pluginmonitor.log                        | Component monitoring plug-in log                                    |
|           | agent_alarm_py.log                       | Log that records alarms upon insufficient NodeAgent file permission |
| okerberos | addRealm.log and modifyKerberosRealm.log | Realm handover log                                                  |
|           | checkservice_detail.log                  | Okerberos health check log                                          |
|           | genKeytab.log                            | keytab generation log                                               |

| Type        | Log File Name                                       | Description                                                                        |
|-------------|-----------------------------------------------------|------------------------------------------------------------------------------------|
|             | KerberosAdmin_genConfigDetail.log                   | Run log file of <b>kadmin.conf</b> generated during start of the kadmin process    |
|             | KerberosServer_genConfigDetail.log                  | Run log file of <b>krb5kdc.conf</b> generated during start of the krb5kdc process  |
|             | oms-kadmind.log                                     | Run log of the kadmin process                                                      |
|             | oms_kerberos_install.log and postinstall_detail.log | Okerberos installation log                                                         |
|             | oms-krb5kdc.log                                     | Run log of the krbkdc process                                                      |
|             | start_detail.log                                    | Okerberos startup log                                                              |
|             | realmDataConfigProcess.log                          | Log file that records the rollback upon a realm handover failure                   |
|             | stop_detail.log                                     | Okerberos stop log                                                                 |
| oldapserver | ldapserver_backup.log                               | Oldapserver backup log                                                             |
|             | ldapserver_chk_service.log                          | Oldapserver health check log                                                       |
|             | ldapserver_install.log                              | Oldapserver installation log                                                       |
|             | ldapserver_start.log                                | Oldapserver startup log                                                            |
|             | ldapserver_status.log                               | Log that records the status of the Oldapserver process                             |
|             | ldapserver_stop.log                                 | Oldapserver stop log                                                               |
|             | ldapserver_wrap.log                                 | Oldapserver service management log                                                 |
|             | ldapserver_uninstall.log                            | Oldapserver uninstallation log                                                     |
|             | restart_service.log                                 | Oldapserver restart log                                                            |
|             | ldapserver_unlockUser.log                           | Log file that records information about unlocking LDAP users and managing accounts |

| Type           | Log File Name           | Description                                                                 |
|----------------|-------------------------|-----------------------------------------------------------------------------|
| metric_agent   | gc.log                  | MetricAgent JVM GC log file                                                 |
|                | metric_agent.log        | Run log file of MetricAgent                                                 |
|                | metric_agent_qps.log    | Log file that records MetricAgent Internal queue length and QPS information |
|                | metric_agent_root.log   | All run log files of MetricAgent                                            |
|                | start.log               | Log file that records information about the MetricAgent startup and stop    |
| omm            | omsconfig.log           | OMS configuration log                                                       |
|                | check_oms_heartbeat.log | OMS heartbeat log                                                           |
|                | monitor.log             | OMS monitoring log                                                          |
|                | ha_monitor.log          | HA_Monitor operation log                                                    |
|                | ha.log                  | HA operation log                                                            |
|                | fms.log                 | Alarm log                                                                   |
|                | fms_ha.log              | HA alarm monitoring log                                                     |
|                | fms_script.log          | Alarm control log                                                           |
|                | config.log              | Alarm configuration log                                                     |
|                | iam.log                 | IAM log                                                                     |
|                | iam_script.log          | IAM control log                                                             |
|                | iam_ha.log              | IAM HA monitoring log                                                       |
|                | config.log              | IAM configuration log                                                       |
|                | operatelog.log          | IAM operation log                                                           |
|                | heartbeatcheck_ha.log   | OMS heartbeat HA monitoring log                                             |
|                | install_oms.log         | OMS installation log                                                        |
| pms_ha.log     | HA monitoring log       |                                                                             |
| pms_script.log | Monitoring control log  |                                                                             |

| Type | Log File Name        | Description                      |
|------|----------------------|----------------------------------|
|      | config.log           | Monitoring configuration log     |
|      | plugin.log           | Monitoring plug-in run log       |
|      | pms.log              | Monitoring log                   |
|      | ha.log               | HA run log                       |
|      | cep_ha.log           | CEP HA monitoring log            |
|      | cep_script.log       | CEP control log                  |
|      | cep.log              | CEP log                          |
|      | config.log           | CEP configuration log            |
|      | omm_gaussdba.log     | GaussDB HA monitoring log        |
|      | gaussdb-<SERIAL>.log | GaussDB run log                  |
|      | gs_ctl-<DATE>.log    | GaussDB control log archive log  |
|      | gs_ctl-current.log   | GaussDB control log              |
|      | gs_guc-current.log   | GaussDB operation log            |
|      | encrypt.log          | Omm encryption log               |
|      | omm_agent_ctl.log    | OMA control log                  |
|      | oma_monitor.log      | OMA monitoring log               |
|      | install_oma.log      | OMA installation log             |
|      | config_oma.log       | OMA configuration log            |
|      | omm_agent.log        | OMA run log                      |
|      | acs.log              | ACS resource log                 |
|      | aos.log              | AOS resource log                 |
|      | controller.log       | Controller resource log          |
|      | floatip.log          | Floating IP address resource log |
|      | ha_ntp.log           | NTP resource log                 |
|      | httpd.log            | Httpd resource log               |
|      | okerberos.log        | Okerberos resource log           |
|      | oldap.log            | OLdap resource log               |



| Type           | Log File Name                                                                | Description                                                       |
|----------------|------------------------------------------------------------------------------|-------------------------------------------------------------------|
|                | tomcat.log                                                                   | Tomcat resource log                                               |
|                | send_alarm.log                                                               | Run log of the HA alarm sending script of the management node     |
|                | feed_watchdog.log                                                            | feed_watchdog resource log                                        |
| timestamp      | restart_stamp                                                                | NodeAgent start time log                                          |
| tomcat         | cas.log and localhost_access_cas_log.log                                     | CAS run log                                                       |
|                | catalina.log, catalina.out, host-manager.log, localhost.log, and manager.log | Tomcat run log                                                    |
|                | localhost_access_web_log.log                                                 | Log that records the access to REST APIs of FusionInsight Manager |
|                | web.log                                                                      | Run log of the web process                                        |
|                | northbound_ftp_sftp.log and snmp.log                                         | Northbound log                                                    |
|                | perfStats.log                                                                | Performance statistics log file                                   |
| watchdog       | watchdog.log and feed_watchdog.log                                           | watchdog run log                                                  |
| update-service | omm_upd_server.log                                                           | UPDServer run log file                                            |
|                | omm_upd_agent.log                                                            | UPDAgent run log file                                             |
|                | update-manager.log                                                           | UPDManager run log file                                           |
|                | install.log                                                                  | Installation log file of the upgrade service                      |
|                | uninstall.log                                                                | Uninstallation log file of the upgrade service                    |

| Type    | Log File Name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Description                         |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
|         | catalina.<Time>.log, catalina.out, host-manager.<Time>.log, localhost.<Time>.log, manager.<Time>.log, manager_access_log.<Time>.txt, web_service_access_log.<Time>.txt, catalina.log, gc-update-service.log.0.current, update-manager.controller, update-web-service.controller, update-web-service.log, commit_rm_distributed.log, commit_rm_upload_package.log, common_omagent_operator.log, forbid_monitor.log, initialize_package_atoms.log, initialize_unzip_pack.log, omm-upd.log, register_patch_pack.log, resume_monitor.logrollback_clear_patch.log, unregister_patch_pack.log, update-rcommupd.log, update-rcupdatemanager.log, and update-service.log | Run log file of the upgrade service |
| upgrade | upgrade.log_<Time>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | OMS upgrade log file                |
|         | rollback.log_<Time>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | OMS rollback log file               |
| sudo    | sudo.log                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Sudo script execution log           |

### Log levels

**Table 7-65** describes the log levels provided by Manager. The priority of log levels, from highest to lowest, is FATAL, ERROR, WARN, INFO, and DEBUG. Logs whose levels are higher than or equal to the set level are printed by the program. The number of printed logs decreases as the set log level increases.

**Table 7-65** Log Levels

| Level | Description                                                                                                                    |
|-------|--------------------------------------------------------------------------------------------------------------------------------|
| FATAL | Logs of this level record fatal error information about the current event handling that may result in a system crash.          |
| ERROR | Logs of this level record error information about the current event handling that indicates a system malfunction.              |
| WARN  | Logs of this level record abnormal information about the current event handling that does not cause the system to malfunction. |
| INFO  | Logs of this level record normal status information about the system and events.                                               |
| DEBUG | Logs of this level record system and its debugging information.                                                                |

**Log format**

Manager log formats are as follows:

**Table 7-66** Log formats

| Type                                                                                   | Component                                                                              | Format                                                                                                                                       | Example                                                                                                                                                                     |
|----------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Controller, HTTPd, Logman, NodeAgent, oKerberos, oldapserver, OMM, Tomcat, and upgrade | Controller, HTTPd, Logman, NodeAgent, oKerberos, oldapserver, OMM, Tomcat, and upgrade | <yyyy-MM-dd HH:mm:ss,SSS> <Log level> <Name of the thread that generates the log> <Message in the log> <Location where the log event occurs> | 2020-06-30 00:37:09,067 INFO [pool-1-thread-1] Completed Discovering Node. com.xxx.hadoop.om.controller.tasks.nodesetup.DiscoverNodeTask.execute(DiscoverNodeTask.java:299) |

## 7.9.2 Viewing MRS Operation Logs

The MRS management console records user operations on MRS clusters and jobs. These logs are used to locate cluster issues, enabling timely resolution.

The MRS console records the following types of operations:

- Cluster operations
  - Create, delete, scale out, and scale in a cluster on the management console.
  - Create a directory and delete a directory and its files on the management console.

- Job operations: Create, stop, and delete jobs on the management console.
- Data operations: IAM user tasks, adding user, and adding user group.

## Viewing MRS Operation Logs

**Step 1** Log in to the MRS console.

**Step 2** In the navigation pane on the left, choose **Operation Logs**.

Operation logs in the log list are sorted in chronological order by default, with the most recent logs displayed at the top.

**Table 7-67** describes various fields in a log.

**Table 7-67** Log description

| Parameter          | Description                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operation Type     | Various types of operations, including: <ul style="list-style-type: none"><li>• Cluster operations</li><li>• Job operations</li><li>• Data operations</li></ul>                                                                                             |
| Operation IP       | IP address where an operation is performed.<br><b>NOTE</b><br>If an MRS cluster fails to be deployed, the cluster is automatically deleted, and the operation logs of the automatically deleted cluster do not contain the <b>Operation IP</b> of the user. |
| Users              | The user who performs the operation                                                                                                                                                                                                                         |
| Operation          | Operation details. The value can contain a maximum of 2048 characters.                                                                                                                                                                                      |
| Time               | Operation time. For a deleted cluster, only logs generated within the last six months are displayed. To view logs generated six months ago, contact Huawei Cloud technical support.                                                                         |
| Enterprise Project | Enterprise project to which the cluster belongs                                                                                                                                                                                                             |

----End

## 7.9.3 Viewing MRS Cluster History

On the MRS management console, you can view terminated or unsubscribed MRS clusters under your account. This allows you to check the status of MRS cluster creation.

### Viewing Information of a Historical Cluster

**Step 1** Log in to the MRS console.

**Step 2** In the navigation pane on the left, choose **Clusters > Cluster History**. On this page, you can view terminated clusters.

Click the name of a terminated cluster. On the displayed page, you can view the cluster configuration, node, auto scaling, component, job, bootstrap action, and tag information.

**Table 7-68** Basic cluster information

| Parameter              | Description                                                                                                                                                  |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Name           | Name of a cluster. The cluster name is set when the cluster is created.                                                                                      |
| Cluster Status         | Cluster status information.                                                                                                                                  |
| Billing Mode           | Billing mode of a cluster. Currently, <b>Pay-per-use</b> and <b>Yearly/Monthly</b> are supported.                                                            |
| Cluster Version        | Cluster version.                                                                                                                                             |
| Cluster Type           | Type of a cluster.                                                                                                                                           |
| Cluster ID             | Unique identifier of a cluster, which is automatically assigned when a cluster is created.                                                                   |
| Created                | Time when a cluster is created.                                                                                                                              |
| Order ID               | Order ID for creating the cluster. This parameter is available only when <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> .                               |
| AZ                     | Availability zone (AZ) in the region of a cluster, which is set when a cluster is created.                                                                   |
| Default Subnet         | Subnet selected during cluster creation.<br>A subnet provides dedicated network resources that are isolated from other networks to enhance network security. |
| VPC                    | VPC selected during cluster creation.<br>A VPC is a secure, isolated, and logical network environment.                                                       |
| OBS Permission Control | Mapping between MRS users and OBS permissions.                                                                                                               |
| Data Connection        | Type of the data connection associated with the cluster.                                                                                                     |
| Agency                 | Agency bound to or modified by the cluster.                                                                                                                  |
| Key Pair               | Name of a key pair, which is set during cluster creation.<br>This parameter is not displayed if password is used as the login mode.                          |

| Parameter               | Description                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Kerberos Authentication | Whether Kerberos authentication is enabled for logging in to Manager.<br><b>NOTE</b><br>Be careful when setting this parameter during the creation of an MRS cluster, as Kerberos authentication cannot be manually enabled or disabled once the cluster is created. To change the Kerberos authentication status, you will need to create a new MRS cluster. |
| Enterprise Project      | Enterprise project to which a cluster belongs. You can only click the name of an enterprise project on the <b>Active Clusters</b> page to access the <b>Enterprise Project Management</b> page.                                                                                                                                                               |
| Security Group          | Security group name of a cluster.                                                                                                                                                                                                                                                                                                                             |
| Streaming Core Node LVM | Whether the Logical Volume Manager (LVM) function is enabled for streaming Core nodes.                                                                                                                                                                                                                                                                        |
| Data Disk Key Name      | Name of the key used to encrypt data disks. The used keys can be managed on the KMS console.                                                                                                                                                                                                                                                                  |
| Data Disk Key ID        | ID of the key used to encrypt data disks.                                                                                                                                                                                                                                                                                                                     |
| Component Version       | Version of each component installed in the cluster.                                                                                                                                                                                                                                                                                                           |
| Agency                  | Binding an agency to the cluster allows the ECSs or BMSs in the cluster to manage specific resources by granting them the necessary permissions.                                                                                                                                                                                                              |

----End

## 7.9.4 Viewing MRS Cluster Audit Logs

The **Audit** page records user operations on Manager. Administrators can view user operation records on Manager. For details, see [Audit logs](#).

This section describes how to view and export audit logs on MRS Manager for post-event tracing, fault cause locating, and responsibility division of security events.

### Viewing Audit Logs (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Audit**. The **Audit** page displays audit information of FusionInsight Manager, including the operation type, risk level, start time, end time, user, host name, service, instance, and operation result.

**Figure 7-27** Audit information list

Audit



| Operation Type | Risk Level | Started              | Completed            | User  | Source name | Host | Service | Instance | Operation Res... |
|----------------|------------|----------------------|----------------------|-------|-------------|------|---------|----------|------------------|
| Lock screen    | Notice     | Aug 5, 2022 16:05... | Aug 5, 2022 16:05... | admin | OMS         | --   | --      | --       | Successful       |
| User login     | Notice     | Aug 5, 2022 15:59... | Aug 5, 2022 15:59... | admin | OMS         | --   | --      | --       | Successful       |
| Unlock screen  | Notice     | Aug 5, 2022 15:59... | Aug 5, 2022 15:59... | admin | OMS         | --   | --      | --       | Successful       |
| Lock screen    | Notice     | Aug 5, 2022 15:55... | Aug 5, 2022 15:55... | admin | OMS         | --   | --      | --       | Successful       |
| User login     | Notice     | Aug 5, 2022 15:55... | Aug 5, 2022 15:55... | admin | OMS         | --   | --      | --       | Successful       |
| User login     | Notice     | Aug 5, 2022 15:55... | Aug 5, 2022 15:55... | admin | OMS         | --   | --      | --       | Successful       |
| User login     | Notice     | Aug 5, 2022 15:54... | Aug 5, 2022 15:54... | admin | OMS         | --   | --      | --       | Successful       |
| User login     | Notice     | Aug 5, 2022 15:45... | Aug 5, 2022 15:45... | admin | OMS         | --   | --      | --       | Successful       |
| User logout    | Notice     | Aug 5, 2022 15:43... | Aug 5, 2022 15:43... | admin | OMS         | --   | --      | --       | Successful       |
| User login     | Notice     | Aug 5, 2022 15:43... | Aug 5, 2022 15:43... | admin | OMS         | --   | --      | --       | Successful       |

10 Total Records: 244 < 1 2 3 4 5 ... 25 >

- You can select audit logs at the **Critical, Major, Minor, or Notice** level from the **All risk levels** drop-down list.
- In **Advanced Search**, you can set filter criteria to query audit logs.
  - You can query audit logs by user management, cluster, service, and health in the **Operation Type** column.
  - In the **Service** column, you can select a service to query corresponding audit logs.

**NOTE**

You can select -- to search for audit logs using all other search criteria except services.

- You can query audit logs by operation result. The options are **All, Successful, Failed, and Unknown**.
- You can click  to manually refresh the current page or click  to choose the columns to display on the page.
  - Click **Export All** to export all audit information at a time, in **TXT** or **CSV** format.

----End

## Viewing Audit Logs (Versions Earlier Than MRS 3.x)

**Step 1** On MRS Manager, click **Audit** and view default audit logs.

**Step 2** If an audit log contains more than 256 characters, click expand button to view full audit details.

- Records are sorted in descending order by the **Occurred** column as the default setting. To change the sorting mode, click **Operation Type, Severity, Occurred, User, Host, Service, Instance, or Operation Result**.
- You can filter all alarms of the same severity, including both cleared and uncleared alarms, by **Severity**.

Exported audit log files contain the following columns:

- Sno**: number of audit log files generated by MRS Manager. The number increases by 1 automatically for each new audit log file generated.

- **Operation Type:** type of a user operation. The options are **Alarm, Auditlog, Backup and Restoration, Cluster, Collect Log, Host, Service, Tenant,** and **User\_Manager**. **User\_Manager** is available only for clusters with Kerberos authentication enabled. Each option contains varying operation types. For example, **Alarm** includes **Export alarms**; **Cluster** includes **Start cluster**; **Tenant** includes **Add tenant**.
- **Severity:** security level of each audit log file, including **Critical, Major, Minor,** and **Informational**.
- **Start Time:** time when an operation starts. The time is CET or CEST.
- **End Time:** time when an operation ends. The time is CET or CEST.
- **User IP Address:** IP address used by a user to perform operations.
- **User:** name of the user who performs operations.
- **Host:** node where user operations are performed. Host information is not recorded in the log file if operations are not performed on any node.
- **Service:** service where user operations are performed. Service information is not recorded in the log file if operations are not performed on any service.
- **Instance:** role instance where user operations are performed. Instance information is not recorded in the log file if operations are not performed on any role instance.
- **Operation Result:** operation result. The options are **Successful, Failed,** and **Unknown**.
- **Content:** execution information of the user operation.

**Step 3** Click **Advanced Search**. In the search area, set search criteria and click **Search** to view audit logs of a specified type. Click **Reset** to clear the search criteria.

 **NOTE**

**Start Time** and **End Time** indicate the start time and end time of a time range. You can search for alarms generated within the time range.

**Step 4** To export all audit log files in the log list, click **Export All**.

**Step 5** To export an audit log file, select it in the log list and click **Export**.

----End

## 7.9.5 Viewing Role Instance Logs of MRS Components

Once an MRS cluster is created, you can easily view the logs of each component role instance and download specific log files on the Manager page. This makes it easier to quickly locate and analyze any issues.

 **NOTE**

The operations described in this section only apply to MRS 3.x or later.

### Viewing Role Instance Logs

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Cluster > Services**, and click a service name. Then click the **Instances** tab of the service and click the name of the target instance to access the instance status page.



**Step 3** In the **Log** area, click the name of a log file to preview its content online.

 **NOTE**

- On the **Hosts** page, click a host name. In the instance list of the host, you can view the log files of all role instances on the host.
- By default, a maximum of 100 lines of logs can be displayed. You can click **Load More** to view more logs. Click **Download** to download the log file to the local PC. For how to download service logs in batches, see [Downloading MRS Cluster Logs](#).

**Figure 7-28** Viewing instance logs

## Log

|                                        |                                          |
|----------------------------------------|------------------------------------------|
| <a href="#">dbservice_audit</a>        | <a href="#">backup</a>                   |
| <a href="#">comonetUserManager</a>     | <a href="#">change_config</a>            |
| <a href="#">checkHaStatus</a>          | <a href="#">cleanupDBService</a>         |
| <a href="#">gaussdbinstall</a>         | <a href="#">gaussdbuninstall</a>         |
| <a href="#">install</a>                | <a href="#">preStartDBService</a>        |
| <a href="#">start_dbserver</a>         | <a href="#">stop_dbserver</a>            |
| <a href="#">dbserver_roll</a>          | <a href="#">dbserver_switchover</a>      |
| <a href="#">status_dbserver</a>        | <a href="#">modifyPassword</a>           |
| <a href="#">modifyDBPwd</a>            | <a href="#">dbservice_metric_collect</a> |
| <a href="#">dbservice_processCheck</a> | <a href="#">dbservice_serviceCheck</a>   |
| <a href="#">ha</a>                     | <a href="#">ha1</a>                      |
| <a href="#">floatip_ha</a>             | <a href="#">gaussDB_ha</a>               |
| <a href="#">ha_monitor</a>             | <a href="#">send_alarm</a>               |
| <a href="#">gaussdb</a>                | <a href="#">gs_guc-current</a>           |
| <a href="#">gs_ctl-current</a>         |                                          |

----End

## 7.9.6 Searching for MRS Cluster Logs Online

On the MRS Manager page, administrators can easily search for and view all component log files by node or component role to identify faults through keyword analysis.

 **NOTE**


The operations described in this section only apply to MRS 3.x or later.

## Searching for Logs Online

**Step 1** Log in to FusionInsight Manager.


**Step 2** Choose **O&M > Log > Online Search**.

**Figure 7-29** Online Search



**Step 3** Set parameters based on [Table 7-69](#). You can select a log search duration or click [✎](#) to specify **Start Time** and **End Time**.

**Table 7-69** Log search parameters

| Parameter        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Search Content   | Keywords or regular expression to be searched for                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Service          | Service or module for which you want to query logs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| File             | Log files to be searched for when only one role is selected                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Lowest Log Level | Lowest level of logs to be queried. After you select a level, the logs of this level and higher levels are displayed.<br>The levels in ascending order are as follows:<br>TRACE < DEBUG < INFO < WARN < ERROR < FATAL                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Host Scope       | <ul style="list-style-type: none"> <li>You can click  to select hosts.</li> <li>Enter the host name of the node for which you want to query logs or the IP address of the management plane.</li> <li>Use commas (,) to separate IP addresses, for example, <b>192.168.10.10,192.168.10.11</b>.</li> <li>Use hyphens (-) to indicate an IP address segment if the IP addresses are consecutive, for example, <b>192.168.10.[10-20]</b>.</li> <li>Use hyphens (-) to indicate an IP address segment if the IP addresses are consecutive, and use commas (,) to separate IP address segments, for example, <b>192.168.10.[10-20,30-40]</b>.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If this parameter is not specified, all hosts are selected by default.</li> <li>A maximum of 10 expressions can be entered at a time.</li> <li>A maximum of 2,000 hosts can be matched for all entered expressions at a time.</li> </ul> |

| Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Advanced Configurations | <ul style="list-style-type: none"><li>• <b>Max Quantity:</b> maximum number of logs that can be displayed at a time. If the number of queried logs exceeds the value of this parameter, the earliest logs will be ignored. If this parameter is not set, the maximum number of logs that can be displayed at a time is not limited.</li><li>• <b>Timeout Duration:</b> log query timeout duration. This parameter is used to limit the maximum log query time on each node. When the query times out, the query is stopped and the logs that have been searched for are still displayed.</li></ul> |

**Step 4** Click **Search**. [Table 7-70](#) describes the fields in search results.

**Table 7-70** Parameters in search results

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                        |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time      | Time when a line of log is generated                                                                                                                                                                                                                                                                                                                               |
| Host Name | Host name of the node where the log file recording the line of log is located                                                                                                                                                                                                                                                                                      |
| Location  | Path of the log file recording the line of log<br>Click the location information to go to the online log browsing page. By default, 100 lines of logs before and 100 lines after the line of log are displayed. You can click <b>Load More</b> on the top or bottom of the page to view more logs. Click <b>Download</b> to download the log file to the local PC. |
| Line No.  | Line number of a line of log in the log file                                                                                                                                                                                                                                                                                                                       |
| Level     | Level of the line of log                                                                                                                                                                                                                                                                                                                                           |
| Log       | Log content                                                                                                                                                                                                                                                                                                                                                        |

 **NOTE**

You can click **Stop** to forcibly stop the search. You can view the search results in the list.

**Step 5** Click **Filter** to filter the logs to display on the page. [Table 7-71](#) lists the fields that you can use to filter logs. After you configure these parameters, click **Filter** to search for logs meeting the search criteria. You can click **Reset** to clear the information that you have filled in.

**Table 7-71** Parameters for filtering logs

| Parameter | Description                             |
|-----------|-----------------------------------------|
| Keywords  | Keywords of the log to be searched for  |
| Host Name | Name of the host to be searched for     |
| Location  | Path of the log file to be searched for |
| Started   | Start time for logs to be searched for  |
| Completed | End time for logs to be searched for    |

----End

## 7.9.7 Downloading MRS Cluster Logs


Manager allows you to export logs from all instances of each service in batches, eliminating the need to manually log in to individual nodes to obtain the logs.

### Downloading MRS Cluster Logs (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Log > Download**.

**Step 3** Select a log download range.

1. **Service:** Click  and select a service.
2. **Host:** Enter IP addresses of the hosts where the service is deployed. You can also click  to select hosts.
3. Click  in the upper right corner and set **Start Time** and **End Time**.

**Step 4** Click **Download**.

The downloaded log package contains the topology information of the start time and end time, helping you quickly find the log you need.

The topology file is named in the format of **topo\_<Topology structure change time>.txt**. The file contains the node IP address, host name, and service instances that reside on the node. (OMS nodes are identified by **Manager:Manager**.)

Example:

```
192.168.204.124|suse-124|
DBService:DBServer;KrbClient:KerberosClient;LdapClient:SlapdClient;LdapServer:SlapdServer;Manager:Manager;meta:meta
```

----End

## Downloading MRS Cluster Logs (Versions Earlier Than MRS 3.x)

- You have obtained the access key ID (AK) and secret access key (SK) of the account.
- A parallel file system has been created in OBS.

**Step 1** On MRS Manager, click **System**.

**Step 2** Click **Export Log** under **Maintenance**.

**Step 3** Select a service for **Service**. Set **Host** to the IP address of the host where the service is deployed. Select the corresponding time for **Start Time** and **End Time**.

**Step 4** In **Export To**, select a path for storing logs. This parameter is available only for clusters with Kerberos authentication enabled.

- **Local PC**: log files are saved to the local environment. Go to [Step 8](#).
- **OBS**: logs are saved to OBS. This is the default option. Go to [Step 5](#).

**Step 5** Set **OBS Path** to the OBS path for storing service log files.

Enter a full path that does not start with a slash (/) and is no more than 900 bytes. The system will automatically create the path if it does not exist.

**Step 6** In **Bucket**, enter the name of the created OBS file system.

**Step 7** Set **AK** and **SK** to the user's access key ID and secret access key.

**Step 8** Click **OK**.

----End

## 7.9.8 Collecting MRS Cluster Service Stack Information

On Manager, administrators can collect stack information of a specified role or instance and download it to a local directory to meet service needs.

The following information can be collected:

1. jstack information.
2. jmap -histo information.
3. jmap -dump information.
4. Thr jstack and jmap-histo information can be collected continuously for comparison.

### NOTE

The operations described in this section only apply to MRS 3.x or later.

## Collecting Stack Information

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Cluster** > **Services** > *Name of the service whose stack information is to be collected*.

**Step 3** On the displayed page, Choose **More** > **Collect Stack Information**.


 **NOTE**

- To collect stack information of multiple instances, go to the instance list, select the desired instances in the instance list and choose **More > Collect Stack Information**.
- To collect stack information of a single instance, click the desired instance and choose **More > Collect Stack Information**.

**Step 4** In the displayed dialog box, select the desired role and content, configure advanced options (retain the default settings if there is no special requirement), and click **OK**.

**Figure 7-30** Collecting stack information




**Collect Stack Information**

 You can collect stack information about all instances of specified roles on this page. If you only need to collect the stack information of some instances, select the instances on the Instance page.

Role:

RegionServer       HMaster

Content:

jstack        jmap -histo        jmap -dump 

Enable continuous collection of jstack and jmap -histo information

Interval:        Duration:


---


Advanced Options

The following options are global policies. Modifying the directory will affect download of previous collected contents.

\* Maximum File Size Printed by jstack and jmap -histo:  MB

\* Number of Archived Files Printed by jstack and jmap -histo:

\* Enable Live Option:        true       false

\* File Directory: 

\* Timeout Period:  s

**Step 5** After the collection is successful, click **Download**.

----End


## Downloading Stack Information

**Step 1** Click **Cluster**, click the name of the desired cluster, click **Services**, and click the target service. Choose **More > Download Stack Information** in the upper right corner.

**Step 2** Select the desired role and content and click **Download** to download the stack information to the local PC.

**Figure 7-31** Downloading stack information

### Download Stack Information

 You can download stack information about all instances of specified roles on this page. If you only need to download the stack information of some instances, select the instances on the Instance page.


Role:

RegionServer       HMaster

Content:

jstack and jmap -histo       jmap -dump

^ Advanced Options

\* File Directory: 


----End

## Clearing Stack Information

- Step 1** Click **Cluster**, click the name of the desired cluster, click **Services**, and click the target service.
- Step 2** Choose **More > Clear Stack Information** in the upper right corner.
- Step 3** Select the desired role and content and configure **File Directory**. Click **OK**.

**Figure 7-32** Clearing stack information


### Clear Stack Information

 To release disk space, you can clear stack information by deleting collection files. You can also stop the continuous collection.


Role:

RegionServer       HMaster

Content:

jstack and jmap -histo       jmap -dump       Continuous collection task 

^ Advanced Options

\* File Directory: 

----End

## 7.9.9 Configuring Default Log Level and Archive File Size for MRS Components

You can change the log levels of FusionInsight Manager. For a specific service, you can change the log level and the log file size to prevent the failure in saving logs due to insufficient disk space.

### NOTE

The operations described in this section only apply to MRS 3.x or later.

### Impact on the System

The services need to be restarted for the new configuration to take effect. During the restart, the services are unavailable.

### Changing the Manager Log Level

1. Log in to the active management node as user **omm**.
2. Run the following command to switch to the required directory:

```
cd ${BIGDATA_HOME}/om-server/om/sbin
```

3. Run the following command to change the log level:

```
./setLogLevel.sh Log level parameters
```

The priorities of log levels are FATAL, ERROR, WARN, INFO, and DEBUG in descending order. Logs whose levels are higher than or equal to the set level are printed. The number of printed logs decreases as the configured log level increases.

- **DEFAULT**: After this parameter is set, the default log level is used.
- **FATAL**: critical error log level. After this parameter is set, only logs of the **FATAL** level are printed.
- **ERROR**: error log level. After this parameter is set, logs of the **ERROR** and **FATAL** levels are printed.
- **WARN**: warning log level. After this parameter is set, logs of the **WARN**, **ERROR**, and **FATAL** levels are printed.
- **INFO** (default): informational log level. After this parameter is set, logs of the **INFO**, **WARN**, **ERROR**, and **FATAL** levels are printed.
- **DEBUG**: debugging log level. After this parameter is set, logs of the **DEBUG**, **INFO**, **WARN**, **ERROR**, and **FATAL** levels are printed.
- **TRACE**: tracing log level. After this parameter is set, logs of the **TRACE**, **DEBUG**, **INFO**, **WARN**, **ERROR**, and **FATAL** levels are printed.

### NOTE

The log levels of components are different from those defined in open-source code.

4. Download and view logs to verify that the log level settings have taken effect. For details, see [Downloading MRS Cluster Logs \(MRS 3.x or Later\)](#).



## Changing the Service Log Level and Log File Size

### NOTE

KrbServer, LdapServer, and DBService do not support the changing of service log levels and log file sizes.

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster > Services**.
- Step 3** Click a service in the service list. On the displayed page, click the **Configuration** page.
- Step 4** On the displayed page, click the **All Configuration** tab. Expand the role instance displayed on the left of the page. Click **Log** of the role to be modified.
- Step 5** Search for each parameter and obtain the parameter description. On the parameter configuration page, select the required log level or change the log file size. The unit of the log file size is MB.

---

### NOTICE

- The system automatically deletes logs based on the configured log size. To save more information, set the log file size to a larger value. To ensure the integrity of log files, you are advised to manually back up the log files to another directory based on the actual service volume before the log files are cleared according to clearance rules.
- Some services do not support change of the log level on the UI.

- 
- Step 6** Click **Save**. In the **Save Configuration** dialog box, click **OK**.
  - Step 7** Download and view logs to verify that the log level settings have taken effect.

----End

## 7.9.10 Configuring the Number of Local Backups of MRS Cluster Audit Logs

Audit logs of cluster components are classified by name and stored in the `/var/log/Bigdata/audit` directory on each cluster node. The OMS automatically backs up the audit log directories at 03:00 every day.

The audit log directory on each node is compressed and named in the `<Node IP address>.tar.gz` format. All compressed files are compressed and named in the `<yyyy-MM-dd_HH-mm-ss>.tar.gz` format and saved in the `/var/log/Bigdata/audit/bk/` directory on the active management node. In addition, the standby management node saves a copy of the file.

By default, a maximum of 90 OMS backup files can be retained. This section describes how to configure the maximum number.

### NOTE

The operations described in this section only apply to MRS 3.x or later.

**Step 1** Log in to the active management node as user **omm**.

 **NOTE**

Perform this operation only on the active management node. This operation is not supported on the standby management nodes; otherwise, the cluster cannot work properly.

**Step 2** Run the following command to switch to the required directory:

```
cd ${BIGDATA_HOME}/om-server/om/sbin
```

**Step 3** Run the following command to change the maximum number of audit log backup files to be retained:

```
./modifyLogConfig.sh -mMaximum number of backup files that can be retained
```

The default value is **90**. The value ranges from **0** to **365**. A larger value means to consume more disk space.

If the following information is displayed, the operation is successful:

```
Modify log config successfully
```

```
----End
```

## 7.9.11 Configuring Dumping for MRS Cluster Audit Logs

The audit logs of FusionInsight Manager are stored in the database by default. If the audit logs are retained for a long time, the disk space of the data directory may be insufficient. To store audit logs to another archive server, administrators can set the required dump parameters to automatically dump these logs. This facilitates the management of audit logs.

If you do not configure the audit log dumping, the system automatically saves the audit logs to a file when the number of audit logs reaches 100,000 pieces. The path is `${BIGDATA_DATA_HOME}/dbdata_om/dumpData/iam/operatelog` on the active management node. The file name format is **OperateLog\_store\_YY\_MM\_DD\_HH\_MM\_SS.csv**. A maximum of 50 historical audit log files can be saved.

 **NOTE**

Archived audit logs will not be displayed on FusionInsight Manager. Only new audit logs generated after the old logs are automatically saved are displayed.

### Configuring Audit Log Dumping to an SFTP Server (MRS 3.x or Later)

The audit logs of FusionInsight Manager are stored in the database by default. If the audit logs are retained for a long time, the disk space of the data directory may be insufficient. To store audit logs to another archive server, administrators can set the required dump parameters to automatically dump these logs. This facilitates the management of audit logs.

If you do not configure the audit log dumping, the system automatically saves the audit logs to a file when the number of audit logs reaches 100,000 pieces. The path is `${BIGDATA_DATA_HOME}/dbdata_om/dumpData/iam/operatelog` on the active management node. The file name format is **OperateLog\_store\_YY\_MM\_DD\_HH\_MM\_SS.csv**. A maximum of 50 historical audit log files can be saved.


 NOTE

Archived audit logs will not be displayed on FusionInsight Manager. Only new audit logs generated after the old logs are automatically saved are displayed.

**Step 1** Log in to FusionInsight Manager.

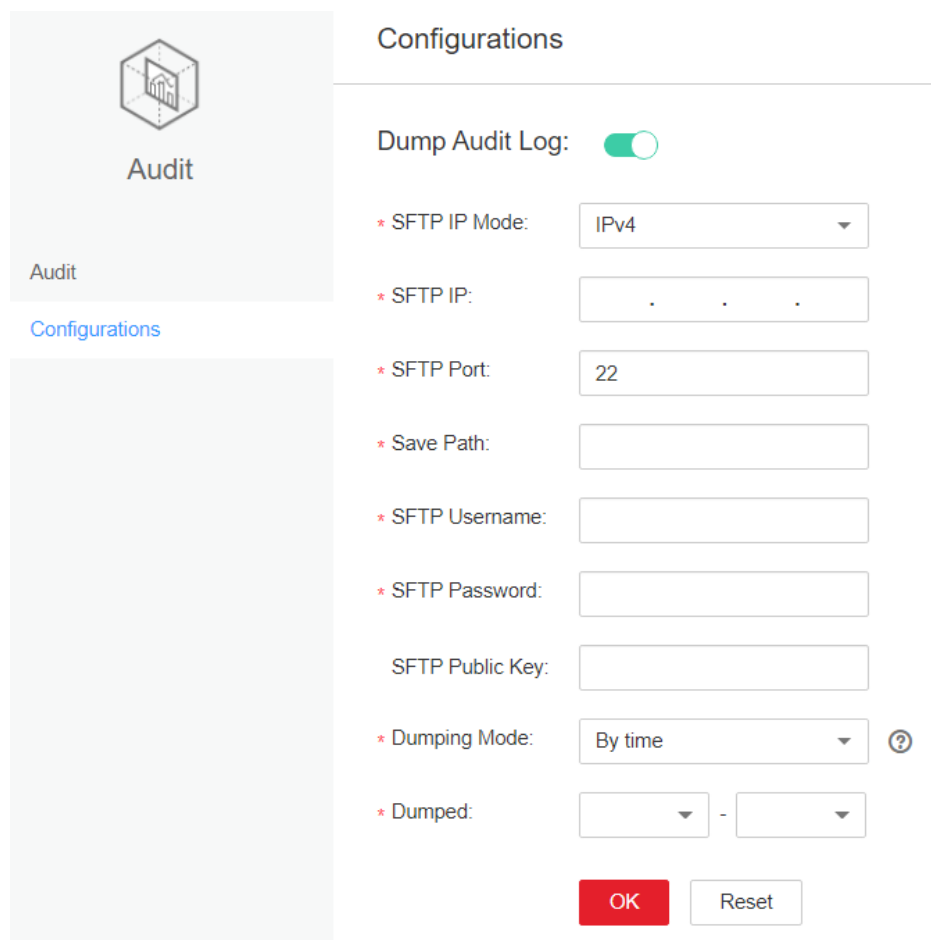
**Step 2** Choose **Audit > Configuration**.

**Step 3** Click the switch on the right of **Audit Log Dumping Flag**.

**Audit Log Dump** is disabled by default. If  is displayed, **Audit Log Dump** is enabled.

**Step 4** Set dumping parameters based on [Table 7-72](#).

**Figure 7-33** Dumping parameters



**Table 7-72** Audit log dump parameters

| Parameter    | Description                                                                       | Value |
|--------------|-----------------------------------------------------------------------------------|-------|
| SFTP IP Mode | Mode of the destination IP address. The value can be <b>IPv4</b> or <b>IPv6</b> . | IPv4  |

| Parameter       | Description                                                                                                                                                                                                                                                                                                                                     | Value                                                                              |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| SFTP IP         | SFTP server for storing dumped audit logs. You are advised to use the SFTP service based on SSH v2 to prevent security risks.                                                                                                                                                                                                                   | <b>192.168.10.51</b><br>(example value)                                            |
| SFTP Port       | Connection port of the SFTP server for storing dumped audit logs                                                                                                                                                                                                                                                                                | <b>22</b><br>(example value)                                                       |
| Save Path       | Path for storing audit logs on the SFTP server                                                                                                                                                                                                                                                                                                  | <b>/opt/omm/oms/auditLog</b><br>(example value)                                    |
| SFTP Username   | User name for logging in to the SFTP server                                                                                                                                                                                                                                                                                                     | <b>root</b><br>(example value)                                                     |
| SFTP Password   | Password for logging in to the SFTP server                                                                                                                                                                                                                                                                                                      | <i>Password for logging in to the SFTP server</i>                                  |
| SFTP Public Key | (Optional) Public key of the SFTP server. Set this parameter to the public key of the SFTP server; or, there may be security risks.                                                                                                                                                                                                             | -                                                                                  |
| Dumping Mode    | Dump mode. Value options are as follows: <ul style="list-style-type: none"> <li>• <b>By Quantity:</b> If the number of pieces of logs reaches the value of this parameter (<b>100000</b> by default), the logs are dumped.</li> <li>• <b>By Time:</b> specifies the date when logs are dumped. The dumping frequency is once a year.</li> </ul> | <ul style="list-style-type: none"> <li>• By Quantity</li> <li>• By Time</li> </ul> |
| Dumping Date    | This parameter is available only when <b>Dumping Mode</b> is set to <b>By time</b> . After you select a dump date, the system starts dumping on this date. The logs to be dumped include all the audit logs generated before January 1 00:00 of the current year.                                                                               | 11-06                                                                              |

 **NOTE**

If the SFTP public key is empty, the system displays a security risk warning. Evaluate the security risk and then save the configuration.

**Step 5** Click **OK** to complete the settings.

 NOTE

Key fields in the audit log dump file are as follows:

- **USERTYPE** indicates the user type. Value **0** indicates a human-machine user, and value **1** indicates a machine-machine user.
- **LOGLEVEL** indicates the security level. Value **0** indicates Critical, value **1** indicates Major, value **2** indicates Minor, and value **3** indicates Warning.
- **OPERATERESULT** indicates the operation result. Value **0** indicates that the operation is successful, and value **1** indicates that the operation failed.

----End

## Configuring Audit Log Dumping to OBS (MRS Versions Earlier Than 3.x)

Storing MRS audit logs in the system may result in insufficient disk space of the data directory. To facilitate audit log management, you can set export parameters to automatically export audit logs to a specified directory on the OBS server.

 NOTE

Audit logs exported to the OBS server include service audit logs and management audit logs.

- Service audit logs are automatically compressed and stored in the `/var/log/Bigdata/audit/bk/` directory on the active management node at 03:00:00 every day. Log files are named in the format of `<yyy-MM-dd_HH-mm-ss>.tar.gz`. By default, only log files generated within the past seven days are stored, while those generated beyond seven days are automatically deleted.
- Management audit logs are exported to OBS starting from the date of the last export until the current execution date of the export task. When there are 100,000 entries in a management audit log file, the system automatically dumps the first 90,000 entries to a local file and keeps the remaining 10,000 in the database. Dumped log files are saved in the `/${BIGDATA_DATA_HOME}/dbdata_om/dumpData/iam/operatelog` directory of the active management node, with a file name format of `OperateLog_store_YY_MM_DD_HH_MM_SS.csv`. A maximum of 50 historical audit log files can be saved.

### Prerequisites

- You have obtained the access key ID (AK) and secret access key (SK) of the account.
- A parallel file system has been created in OBS.

### Procedure

**Step 1** On MRS Manager, choose **System**.

**Step 2** Choose **Export Audit Log** under **Maintenance**.

**Table 7-73** Parameters for exporting audit logs

| Parameter        | Value                                                          | Description                                                                                                                                                                                                |
|------------------|----------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Export Audit Log | <ul style="list-style-type: none"><li>On</li><li>Off</li></ul> | (Mandatory) Whether to enable audit log export. The options are: <ul style="list-style-type: none"><li><b>On</b>: Audit log export is enabled.</li><li><b>Off</b>: Audit log export is disabled.</li></ul> |
| Started          | 07/24/2017 09:00:00<br>(example)                               | (Mandatory) Start time for exporting audit logs.                                                                                                                                                           |
| Period           | 1 day (example)                                                | (Mandatory) Interval for exporting audit logs. The interval ranges from 1 to 5 days.                                                                                                                       |
| Bucket           | mrs-bucket (example)                                           | (Mandatory) Name of the OBS file system to which audit logs are exported.                                                                                                                                  |
| OBS Path         | /opt/omm/oms/<br>auditLog (example)                            | (Mandatory) OBS path to which audit logs are exported.                                                                                                                                                     |
| AK               | XXX (example)                                                  | (Mandatory) User's access key ID.                                                                                                                                                                          |
| SK               | XXX (example)                                                  | (Mandatory) User's secret access key.                                                                                                                                                                      |

**NOTE**

Service and management audit log files are stored in **service\_auditlog** and **manager\_auditlog** on OBS, respectively.

----End

## 7.10 MRS Cluster Security Configuration

### 7.10.1 Cluster Mutual Trust Management

#### 7.10.1.1 Overview of Mutual Trust Between MRS Clusters

##### Function Description

By default, users of a big data cluster in security mode can only access resources in the cluster but cannot perform identity authentication or access resources in other clusters in security mode.

## Concepts of Cluster Mutual Trust

- **Domain**

The secure usage scope of users in each system is called a domain. Each FusionInsight Manager must have a unique domain name. Cross-Manager access allows users to use resources across domains.
- **User Encryption**

Mutual trust can be configured across FusionInsight Managers. The current Kerberos server supports only the aes256-cts-hmac-sha1-96:normal and aes128-cts-hmac-sha1-96:normal encryption types for encrypting cross-domain users, and the encryption types cannot be changed.
- **User Authentication**

After cross-Manager mutual trust is configured, if a user with the same name exists in two systems and the user in the peer system has the permission to access a resource in that system, this user can also access the remote resource.
- **Direct Mutual Trust**

The system saves the mutual trust ticket of the peer system in two clusters with mutual trust configured and uses the mutual trust ticket to access the peer system.

### 7.10.1.2 Changing the System Domain Name of an MRS Cluster

#### Scenario

The secure usage scope of users in each system is called a domain. Each system must have a unique domain name. The domain name of FusionInsight Manager is generated during installation. The system administrator can change the domain name on FusionInsight Manager.

---

#### NOTICE

- Changing the system domain name is a high-risk operation. Before performing operations in this section, ensure that the OMS data has been backed up by referring to [Backing Up Manager Data \(MRS 3.x and Later Versions\)](#).
  - This topic is available for MRS 3.1.2 or later.
- 

#### Impact on the System

- During the configuration, all of the clusters need to be restarted and are unavailable during restart.
- After the domain name is changed, the passwords of the Kerberos administrator and OMS Kerberos administrator will be initialized. You need to use the default passwords and then change the passwords. If a component user whose password is generated randomly by the system is used for identity authentication, see [Downloading MRS Cluster User Credentials](#) to download the keytab file again.
- After the domain name is changed, passwords of the **admin** user, component user, and human-machine user added by the system administrator before the

domain name change will be reset to the same one. Change these passwords. The reset password consists of two parts: one part is generated by the system and the other is set by the user. The system generating part is **Admin@123**, which is the default password. For details about the user-defined part, see descriptions of **Password Suffix** in [Table 7-75](#). For example, if the system generates **Admin@123** and the user sets **Test#\$%@123**, the new password after reset is **Admin@123Test#\$%@123**.

- The new password must meet the password policies. To obtain the new human-machine user password, log in to the active OMS as user **omm** and run the following script:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/get_reset_pwd.sh Password
suffix user_name
```

- *Password suffix* is a parameter set by the user. If it is not specified, the default value **Admin@123** is used.
- *user\_name* is optional. The default value is **admin**.
- There can be security risks if a command contains the authentication password. You are advised to disable the command recording function (history) before running the command.

Example:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/get_reset_pwd.sh Test#$
%@123
```

To get the reset password after changing cluster domain name.

```
pwd_min_len : 8
pwd_char_types : 4
```

The password reset after changing cluster domain name is: "Admin@123Test#\$%@123"

In this example, **pwd\_min\_len** and **pwd\_char\_types** indicate the minimum password length and number of password character types respectively defined in the password policies. **Admin@123Test#\$%@123** indicates the human-machine user password after the system domain name is changed.

- After the system domain name is changed, the reset password consists of two parts: one part is generated by the system and the other is set by the user. The reset password must meet the password policies. If the password is not long enough, one or multiple at signs (@) are added between **Admin@123** and the user-defined part. If there are five character types, a space is added after **Admin@123**.

When the user-defined part is **Test@123** and the default user password policy is used, the new password is **Admin@123Test@123**. The password contains 17 characters of four types. To meet the current password policy, the new password is processed according to [Table 7-74](#).

**Table 7-74** Password processing

| Minimum Password Length | Number of Character Types | Processing Against the Password Policy | New Password       |
|-------------------------|---------------------------|----------------------------------------|--------------------|
| 8 to 17 characters      | 4                         | The user password policy is met.       | Admin@123Test@123  |
| 18 characters           | 4                         | Add an at sign (@).                    | Admin@123@Test@123 |



| Minimum Password Length | Number of Character Types | Processing Against the Password Policy | New Password            |
|-------------------------|---------------------------|----------------------------------------|-------------------------|
| 19 characters           | 4                         | Add two at signs (@).                  | Admin@123@@Test@123     |
| 8 to 18 characters      | 5                         | Add a space.                           | Admin@123<br>Test@123   |
| 19 characters           | 5                         | Add a space and an at sign (@).        | Admin@123<br>@Test@123  |
| 20 characters           | 5                         | Add a space and two at signs (@).      | Admin@123<br>@@Test@123 |

- After the system domain name is changed, download the **keytab** file for the machine-machine user added by the system administrator before the domain name is changed.
- After the system domain name is changed, download and install the client again.
- After the system domain name is changed, if there is any running HetuEngine compute instance, restart the instance.

## Prerequisites


- The system administrator has clarified service requirements and planned domain names for the systems.  
A domain name can contain only uppercase letters, numbers, periods (.), and underscores (\_), and must start with a letter or number, for example, **DOMAINA.HW** and **DOMAINB.HW**.
- The running status of all components in the Manager clusters is **Normal**.
- The **acl.compare.shortName** parameter of the ZooKeeper service of all clusters in Manager is set to default value **true**. Otherwise, change the value to **true** and restart the ZooKeeper service.

## Changing the System Domain Name

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **System > Permission > Domain and Mutual Trust**.

**Figure 7-34** Domain and mutual trust



**Step 3** Modify required parameters.

**Table 7-75** Related parameters

| Parameter       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Domain    | Planned domain name of the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Password Suffix | Part of the password set by the user after the password of the human-machine user is reset. This parameter is mandatory. The default value is <b>Admin@123</b> .<br><b>NOTE</b><br>This parameter takes effect only after <b>Local Domain</b> is modified. The following conditions must be met: <ul style="list-style-type: none"> <li>The password ranges from 8 to 16 characters.</li> <li>The password must contain at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (^~!@#\$\$%^&amp;*()-_+= [{}];:','&lt;.&gt;/? and spaces).</li> </ul> |

**Step 4** Determine whether to set up mutual trust between managers across cluster.

- If you need to do so, perform operations by referring to [Configuring Mutual Trust Between MRS Clusters](#) and skip subsequent steps in this section.
- If you do not need such configuration, go to [Step 5](#).

**Step 5** Click **OK**. Proceed with the subsequent steps only after the modification is complete.

**Step 6** Log in to the active management node as user **omm**.

**Step 7** Run the following command to update the domain configuration:


```
sh ${BIGDATA_HOME}/om-server/om/sbin/restart-RealmConfig.sh
```

The command is executed successfully if the following information is displayed:

Modify realm successfully. Use the new password to log in to FusionInsight again.

 NOTE

After the restart, some hosts and services cannot be accessed and an alarm is generated. This problem can be automatically resolved in about 1 minute after **restart-RealmConfig.sh** is run.

**Step 8** Log in to FusionInsight Manager using the new password of user **admin** (for example, **Admin@123Admin@123**). On the dashboard, click  or **More** and select **Restart**.

In the displayed dialog box, enter the password of the current login user and click **OK**.

In the displayed dialog box, click **OK**. Wait for a while until a message indicating that the operation is successful is displayed. Click **Finish**.

**Step 9** Log out of FusionInsight Manager and then log in again. If the login is successful, the configuration is successful.

**Step 10** Log in to the active management node as user **omm** and run the following command to update the configurations of the job submission client:

```
sh /opt/executor/bin/refresh-client-config.sh
```

**Step 11** If a HetuEngine compute instance is running, restart the compute instance.

1. Log in to FusionInsight Manager as the user who is used to access the HetuEngine web UI.
2. Choose **Cluster > Services > HetuEngine** to go to the HetuEngine service page.
3. In the **Basic Information** area on the **Dashboard** page, click the link next to **HSConsole WebUI**. The HSConsole page is displayed.
4. For a running compute instance, click **Stop** in the **Operation** column. After the compute instance is in the **Stopped** state, click **Start** to restart the compute instance.

----End

### 7.10.1.3 Configuring Mutual Trust Between MRS Clusters

When clusters in different Manager systems with security modes need to access each other's resources, system administrators can establish mutual trust between the systems, allowing external system users to access resources in the local system.

Without configuring cross-cluster trust, resources in each cluster can only be accessed by local users. The scope of secure usage for each system user is defined as a domain, and unique domain names must be defined for different Manager systems. Cross-Manager access is essentially cross-domain usage by users.

 NOTE

In MRS 3.x or later, a maximum of 500 mutual-trust clusters can be configured.

## Impact on the System

- Once cross-Manager cluster mutual trust is configured, users from an external system can be utilized in the local system. The system administrator should regularly review user permissions in Manager to ensure they align with enterprise service and security requirements.
- If you set up mutual trust between clusters, affected services need to be restarted and services will be interrupted.
- After cross-Manager cluster mutual trust is configured, internal Kerberos users **krbtgt/Local cluster domain name@External cluster domain name** and **krbtgt/External cluster domain name@Local cluster domain name** are added to the two mutually trusted clusters. The internal users cannot be deleted.
  - For MRS 2.x or earlier, the default password is **Crossrealm@123**.
  - For MRS 3.x or later, the system administrator should regularly update the passwords of the four users in the mutually trusted systems and ensure they are consistent. For details, see [Changing the Passwords for MRS Cluster Component Running Users](#). When the passwords are changed, the connectivity between cross-cluster service applications may be affected.
- For MRS 3.x or later, if the system domain name is changed and there is any running HetuEngine compute instance, restart the compute instance.
- For MRS 3.x or later, after cross-Manager cluster mutual trust is configured, the clients of each cluster need to be redownloaded and reinstalled.
- After cross-Manager cluster mutual trust is configured, you need to check whether the system works properly and how to access resources of the peer system as a user of the local system. For details, see [Configuring User Permissions for Mutually Trusted MRS Clusters](#).

## Prerequisites

- The system administrator has clarified service requirements and planned domain names for the systems. A domain name can contain uppercase letters, numbers, periods (.), and underscores (\_), and must start with a letter or number. For example, **DOMAINA.HW** and **DOMAINB.HW**.
- The domain names of the two Managers are different. When an ECS or BMS cluster is created on MRS, a unique system domain name is randomly generated. Generally, you do not need to change the system domain name.
- The two clusters do not have the same host name or the same IP address.
- The system time of the two clusters is consistent, and the NTP services in the two systems use the same clock source.
- The running status of all components in the Manager clusters is **Normal**.
- The two clusters are in the same VPC. If they are not, create a VPC peering connection between them. For details, see [VPC Peering Connection](#).
- For MRS 3.x or later, the **acl.compare.shortName** parameter for the ZooKeeper service in all clusters within Manager should be set to the default value of **true**. Otherwise, change the value to **true** and restart the ZooKeeper service.

## Configuring Mutual Trust Between MRS Clusters (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager of one of the two clusters.



**Step 2** Choose **System > Permission > Domain and Mutual Trust**.

**Step 3** Modify **Peer Mutual Trust Domain**.

**Table 7-76** Parameters

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| realm_name | Enter the domain name of the peer system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| ip_port    | <p>Enter the KDC address of the peer system.</p> <p>Value format: <i>IP address of the node accommodating the Kerberos service in the peer system:Port number</i></p> <ul style="list-style-type: none"><li>• In dual-plane networking, enter the service plane IP address.</li><li>• If an IPv6 address is used, the IP address must be enclosed in square brackets ([ ]).</li><li>• Use commas (,) to separate the KDC addresses if the active and standby Kerberos services are deployed or multiple clusters in the peer system need to establish mutual trust with the local system.</li><li>• You can obtain the port number from the <b>kdc_ports</b> parameter of the KrbServer service. The default value is <b>21732</b>. To obtain the IP address of the node where the service is deployed, click the <b>Instance</b> tab on the KrbServer page and view <b>Service IP Address</b> of the KerberosServer role.</li></ul> <p>For example, if the Kerberos service is deployed on nodes at <b>10.0.0.1</b> and <b>10.0.0.2</b> that have established mutual trust with the local system, the parameter value is <b>10.0.0.1:21732,10.0.0.2:21732</b>.</p> |

 **NOTE**

If you need to configure mutual trust for multiple Managers, click  to add a new item and set parameters. Click  to delete unnecessary configurations.

**Step 4** Click **OK**.

**Step 5** Log in to the active management node as user **omm** and run the following command to update the domain configuration:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/restart-RealmConfig.sh
```

The command is executed successfully if the following information is displayed:

Modify realm successfully. Use the new password to log in to FusionInsight again.

After the restart, some hosts and services cannot be accessed and an alarm is generated. This problem can be automatically resolved in about 1 minute after **restart-RealmConfig.sh** is run.

**Step 6** Log in to FusionInsight Manager and restart the cluster or configure expired instances:

Check whether the system domain name of Manager is changed.

- If the system domain name is changed, click **\*\*\*** or **More** on the home page, click **Start**, enter the password, select the checkbox for confirming the impact, and click **OK**. Wait until the cluster is started successfully.
- If the system domain name is not changed, click **\*\*\*** or **More** on the home page and select **Restart Configuration-Expired Instances**. Enter the password, select the checkbox for confirming the impact, and click **OK**. Wait until the service is restarted.

**Step 7** Log out of FusionInsight Manager and then log in again. If the login is successful, the configuration is successful.

**Step 8** If a HetuEngine compute instance is running, restart the compute instance.

1. Log in to FusionInsight Manager as the user who is used to access the HetuEngine web UI.
2. Choose **Cluster > Services > HetuEngine** to go to the HetuEngine service page.
3. In the **Basic Information** area on the **Dashboard** page, click the link next to **HSConsole WebUI**. The HSConsole page is displayed.
4. For a running compute instance, click **Stop** in the **Operation** column. After the compute instance is in the **Stopped** state, click **Start** to restart the compute instance.

**Step 9** Log in to FusionInsight Manager of another cluster and repeat the preceding steps.

----End

## Configuring Mutual Trust Between MRS Clusters (MRS 2.x or Earlier)

**Step 1** On the MRS management console, query all security groups of the two clusters.

- If the security groups of the two clusters are the same, go to [Step 3](#).
- If the security groups of the two clusters are different, go to [Step 2](#).

**Step 2** On the VPC management console, choose **Access Control > Security Groups**. On the **Security Groups** page, locate the row containing the target security group, click **Manage Rule** in the **Operation** column.

On the **Inbound Rules** tab page, click **Add Rule**. In the **Add Inbound Rule** dialog box that is displayed, configure related parameters.

- **Priority:** The value ranges from 1 to 100. The default value is **1**, which indicates the highest priority. A smaller value indicates a higher priority.
- **Action:** Select **Allow**.
- **Protocol & Port:** Choose **Protocols > All**.

- **Type:** Select **IPv4** or **IPv6**.
- **Source:** Select **Security group** and the security group of the peer cluster.
  - To add an inbound rule to the security group of cluster A, set **Source** to **Security group** and the security group of cluster B (peer cluster).
  - To add an inbound rule to the security group of cluster B, set **Source** to **Security group** and the security group of cluster A (peer cluster).

 **NOTE**

For a common cluster with Kerberos authentication disabled, perform step [Step 1](#) to [Step 2](#) to configure cross-cluster mutual trust. For a security cluster with Kerberos authentication enabled, after completing the preceding steps, proceed to the following steps for configuration.

**Step 3** Log in to MRS Manager of the two clusters separately. Click **Service** and check whether the **Health Status** of all components is **Good**.

- If yes, go to [Step 4](#).
- If no, contact O&M personnel for troubleshooting.

**Step 4** Query configuration information.



1. On MRS Manager of the two clusters, choose **Services > KrbServer > Instance**. Query the **OM IP Address** of the two KerberosServer hosts.
2. Click **Service Configuration**. Set **Type** to **All**. Choose **KerberosServer > Port** in the navigation tree on the left. Query the value of **kdc\_ports**. The default value is **21732**.
3. Click **Realm** and query the value of **default\_realm**.

**Step 5** On MRS Manager of either cluster, modify the **peer\_realms** parameter.

**Table 7-77** Parameter description

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| realm_name | Domain name of the mutual-trust cluster, that is, the value of <b>default_realm</b> obtained in step <a href="#">4</a> .                                                                                                                                                                                                                                               |
| ip_port    | KDC address of the peer cluster. Format: <i>IP address of a KerberosServer node in the peer cluster:kdc_port</i><br>The addresses of the two KerberosServer nodes are separated by a comma. For example, if the IP addresses of the KerberosServer nodes are 10.0.0.1 and 10.0.0.2 respectively, the value of this parameter is <b>10.0.0.1:21732,10.0.0.2:21732</b> . |

 **NOTE**

- To deploy trust relationships with multiple clusters, click  to add items and specify relevant parameters. To delete an item, click .
- A cluster can have trust relationships with a maximum of 16 clusters. By default, no trust relationship exists between different clusters that are trusted by a local cluster.

- Step 6** Click **Save Configuration**. In the dialog box that is displayed, select **Restart the affected services or instances** and click **OK**. If you do not select **Restart the affected services or instances**, manually restart the affected services or instances.
- After **Operation successful** is displayed, click **Finish**.
- Step 7** Exit MRS Manager and log in to it again. If the login is successful, the configurations are valid.
- Step 8** Log in to MRS Manager of the other cluster and repeat step [Step 5](#) to [Step 7](#).
- Step 9** Perform subsequent operations by referring to [Updating the Client Configuration of Mutually Trusted Clusters \(MRS 2.x or Earlier\)](#).
- End

## Updating the Client Configuration of Mutually Trusted Clusters (MRS 2.x or Earlier)

After cross-cluster mutual trust is configured, the service configuration parameters are modified on MRS Manager and the service is restarted. Therefore, you need to prepare the client configuration file again and update the client.

Scenario 1:

Cluster A and cluster B (peer and mutually trusted clusters) are of the same type, for example, analysis cluster or streaming cluster. Refer to [Updating Client Configurations \(Version 2.x or Earlier\)](#) to update the client configuration files accordingly.

- Update the client configuration file of cluster A.
- Update the client configuration file of cluster B.

Scenario 2:

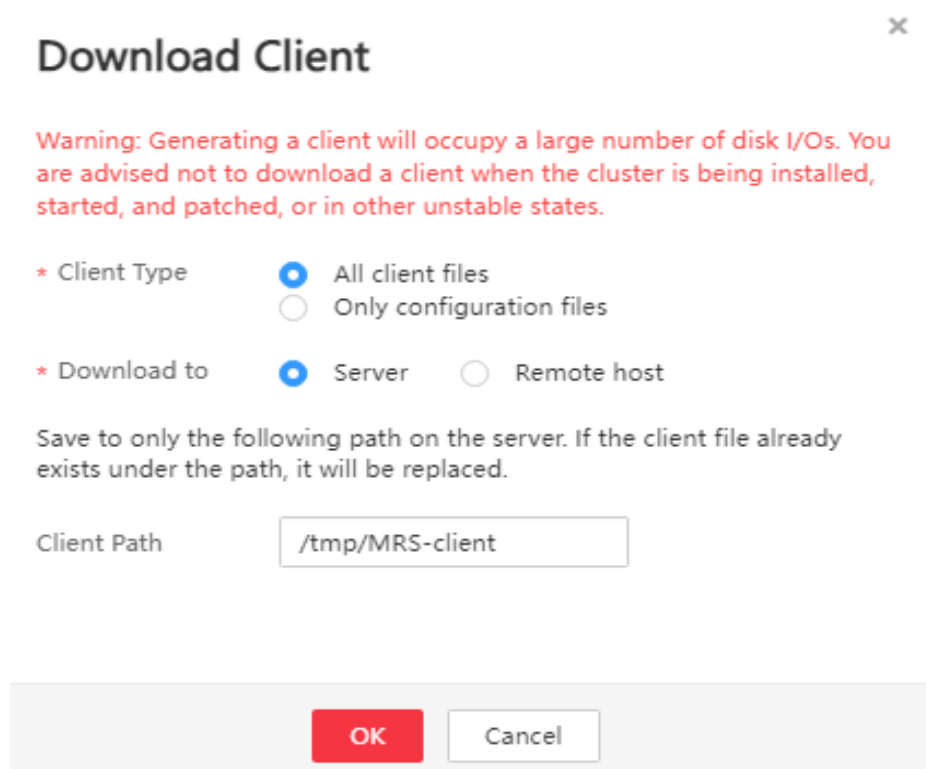
Cluster A and cluster B (peer cluster and mutually trusted cluster) are the different type. Perform the following steps to update the configuration files.

- Update the client configuration file of cluster A to cluster B.
- Update the client configuration file of cluster B to cluster A.
- Update the client configuration file of cluster A.
- Update the client configuration file of cluster B.

- Step 1** Log in to MRS Manager of cluster A.
- Step 2** Click **Services**, and then **Download Client**.



Figure 7-35 Downloading a client



**Step 3** Set **Client Type** to **Only configuration files**.

**Step 4** Set **Download to** to **Remote host**.

**Step 5** Set **Host IP Address** to the IP address of the active Master node of cluster B, **Host Port** to 22, and **Save Path** to **/tmp**.

- If the default port **22** for logging in to cluster B using SSH is changed, set **Host Port** to a new port.
- The value of **Save Path** contains a maximum of 256 characters.

**Step 6** Set **Login User** to **root**.

If another user is used, ensure that the user has permissions to read, write, and execute the save path.

**Step 7** Select **Password** or **SSH Private Key** for **Login Mode**.

- **Password:** Enter the password of user **root** set during cluster creation.
- **SSH Private Key:** Select and upload the key file used for creating the cluster.

**Step 8** Click **OK** to generate a client file.

- If the following information is displayed, the client package is saved. Click **Close**.  
Client files downloaded to the remote host successfully.
- If the following information is displayed, check the username, password, and security group configurations of the remote host. Ensure that the username and password are correct and an inbound rule of the SSH (22) port has been added to the security group of the remote host. And then, go to **Step 2** to download the client again.  
Failed to connect to the server. Please check the network connection or parameter settings.

**Step 9** Log in to the ECS of cluster B using VNC. For details, see [Logging In to a Windows ECS Using VNC](#).

All images support Cloud-Init. The preset username for Cloud-Init is **root**, and the password is the one set during cluster creation.

**Step 10** Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

```
cd /opt/Bigdata/client
```

**Step 11** Run the following command to update the client configuration of cluster A to cluster B:

```
sh refreshConfig.sh Client installation directory Full path of the client configuration file package
```

For example, run the following command:

```
sh refreshConfig.sh /opt/Bigdata/client /tmp/MRS_Services_Client.tar
```

If the following information is displayed, the configurations have been updated successfully.

```
ReFresh components client config is complete.
Succeed to refresh components client config.
```

 **NOTE**

You can also refer to method 2 in [Updating Client Configurations \(Version 2.x or Earlier\)](#) to perform operations in [Step 1](#) to [Step 11](#).

**Step 12** Repeat step [Step 1](#) to [Step 11](#) to update the client configuration file of cluster B to cluster A.

**Step 13** Refer to [Updating Client Configurations \(Version 2.x or Earlier\)](#) to update the client configuration file of the local cluster.

- Update the client configuration file of cluster A.
- Update the client configuration file of cluster B.

----End

### 7.10.1.4 Configuring User Permissions for Mutually Trusted MRS Clusters


Once cross-Manager cluster mutual trust is configured, grant user access permissions on FusionInsight Managers to enable them to perform service operations in the mutually trusted Managers.

#### Prerequisites


The mutual trust relationship has been configured between two clusters (clusters A and B). The clients of the clusters have been updated.

### Configuring User Permissions for Mutually Trusted Clusters (MRS 3.x or Later)

**Step 1** Log in to the local FusionInsight Manager.

- Step 2** Choose **System > Permission > User** to check whether the target user exists.
- If yes, go to [Step 3](#).
  - If no, go to [Step 4](#).
- Step 3** Click  on the left of the target user, and check whether the permissions assigned to the user group of the user and the roles meet service requirements. If not, create a role and bind the role to the user, or modify the user group or role permissions of the user.
- Step 4** Create a user required by the service operations and associate the required user group or role. For details, see [Creating a User \(MRS 3.x and Later\)](#).
- Step 5** Log in to the other FusionInsight Manager and repeat [Step 2](#) to [Step 4](#) to create a user with the same name and set permissions.
- End

## Configuring User Permissions for Mutually Trusted Clusters (MRS 2.x or Earlier)

- Step 1** Log in to MRS Manager of cluster A and choose **System > Manage User**. Check whether cluster A has accounts that are the same as those of cluster B.
- If yes, go to [Step 2](#).
  - If no, go to [Step 3](#).
- Step 2** Click  on the left side of the username to view detailed user information. Verify if the user's group and role meet the service requirements.
- For example, user **admin** of cluster A has the permission to access and create files in the **/tmp** directory of cluster A. Then go to [Step 4](#).
- Step 3** Create the accounts in cluster A and bind the accounts to the user group and roles required by the services. Then go to [Step 4](#).
- Step 4** Choose **Service > HDFS > Instance**. Query the **OM IP Address of NameNode (Active)**.
- Step 5** Log in to the client of cluster B.
- For example, if you updated the client on the Master2 node, you must log in to the Master2 node to use the client.
- Step 6** Run the following command to access the **/tmp** directory of cluster A.
- ```
hdfs dfs -ls hdfs://192.168.6.159:9820/tmp
```
- In the preceding command, **192.168.6.159** is the IP address of the active NameNode of cluster A; **9820** is the default port for communication between the client and the NameNode.
- Step 7** Run the following command to create a file in the **/tmp** directory of cluster A:
- ```
hdfs dfs -touchz hdfs://192.168.6.159:9820/tmp/mrstest.txt
```
- If you can find the **mrstest.txt** file in the **/tmp** directory of cluster A, the cross-cluster mutual trust is successfully configured.
- End

## 7.10.2 Replacing MRS Cluster Certificates

### 7.10.2.1 Replacing the CA Certificate

#### Scenario

The MRS CA certificate is used for data encryption during the communication between the client and the server of a component to ensure communication security. You can replace the CA certificate on FusionInsight Manager to ensure product security. This operation is applicable to the following scenarios:

- After the cluster is installed for the first time, import an enterprise certificate.
- If the enterprise certificate has expired or security hardening is required, replace it with a new certificate.

After the CA certificate is replaced, the certificates that are used by HDFS, YARN, MapReduce, HBase, Loader, Hue, Flink (MRS 3.2.0 or later) Oozie, Hive, Tomcat, CAS, HTTPD, and LDAP in MRS will be automatically updated.

The certificate file and key file can be applied for from the enterprise certificate center or generated by the cluster user.

#### NOTE

- Only CA certificates that can be issued and in **X.509** format can be imported in FusionInsight.
- FusionInsight requires that the OS encoding format be **en\_US.UTF-8** or **POSIX**. Otherwise, the certificate function will be abnormal.
- If an isolated faulty node exists in the current cluster, the CA certificate of the node will not be replaced. After the node is de-isolated, you need to reinstall the services running on the node to ensure that the node and the cluster use the same CA certificate.
- This topic is available for MRS 3.x or later.

#### Impact on the System

- The MRS system must be restarted during the replacement and cannot be accessed or provide services.
- After the certificate is replaced, the certificates used by all components and FusionInsight Manager modules are automatically updated.
- After the certificate is replaced, you need to reinstall the certificate in the local environment where the certificate is not trusted.

#### Prerequisites

- You have obtained the files to be imported to the MRS cluster, including the CA certificate file (**\*.crt**), key file (**\*.key**), and file that saves the key file password (**password.property**). The certificate name and key name support letters and digits.
- You have prepared a password for accessing the key file, for example, **Userpwd@123**.

To avoid potential security risks, the password must meet the following complexity requirements:

- Contains at least 8 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!?,,:-\_'(){}[]/<>@#\$\$%^&\*+|\=).
- When applying for a certificate from the certificate center, provide the password for accessing the key file and apply for the certificate files in CRT, CER, CERT, and PEM formats and the key files in KEY and PEM formats. The applied certificate must have the issuing function.

## Replacing the CA Certificate

**Step 1** Log in to any management node in the cluster as user **omm**.

**Step 2** Select a method for generating certificate files and key files.

- If the certificate is generated by the certificate center, save the certificate file and key file to the **omm** user directory on the management node.

### NOTE

If the obtained certificate file is not in the **.crt** format and the key file is not in the **.key** format, run the following commands to change the file formats:

```
mv Certificate name.Certificate formatCertificate name.crt
```

```
mv Key name.Key format Key name.key
```

For example, run the following commands to name the certificate file **ca.crt** and name the key file **ca.key**:

```
mv server.cer ca.crt
```

```
mv server_key.pem ca.key
```

- If the certificate is generated by the cluster user, run the following commands to generate the certificate file and key file in the **omm** user directory on the management node:

a. Generate the key file.

Run the following command to check whether the OpenSSL version is 1.1.1 or later:

```
/usr/bin/openssl version
```

- If yes, run the following command:

```
openssl genrsa -out Key name.key -aes256 3072
```

- If no, run the following command:

```
openssl genrsa -out Key name.key -aes256 3072 -sha256
```

For example, to generate the key file **ca.key**, run the following command:

```
openssl genrsa -out ca.key -aes256 3072 -sha256
```

Enter the password twice as prompted, and press **Enter**.

```
Enter pass phrase for ca.key:
```

```
Verifying - Enter pass phrase for ca.key:
```

b. Generate the certificate file.

```
openssl req -new -x509 -days 1825 -key Key name.key -out Certificate name.crt -subj "/C=cn/ST=guangdong/L=shenzhen/O=huawei/OU=huawei/CN=huawei" -sha256
```

For example, to generate the certificate file **ca.crt**, run the following command:

```
openssl req -new -x509 -days 1825 -key ca.key -out ca.crt -subj "/
C=cn/ST=guangdong/L=shenzhen/O=huawei/OU=huawei/
CN=huawei" -sha256
```

Enter the password for the key file as prompted, and press **Enter**.

Enter pass phrase for ca.key:

**Step 3** Run the following command in the **omm** user directory on the management node to save the password for accessing the key file.

```
sh ${BIGDATA_HOME}/om-server/om/sbin/genPwFile.sh
```

Enter the password twice as prompted, and press **Enter**. After being encrypted, the password is saved in **password.property**.

Please input key password:

Please Confirm password:

#### NOTE

- The **password.property** file generated on the node you have logged is available only for the current cluster and cannot be used for other clusters. The file contains security information. Keep it secure and control the access permission.
- In active/standby DR scenarios, the **genPwFile.sh** script must be executed on both the active and DR cluster nodes, and the same password must be entered for the two clusters.

**Step 4** Compress the three files in the **.tar** format and save them to the local computer.

```
tar -cvf Package name Certificate name .crt Key name .key password.property
```

For example, **tar -cvf test.tar ca.crt ca.key password.property**

#### NOTE

In active/standby DR scenarios, run this command on each cluster node.

**Step 5** Log in to FusionInsight Manager and choose **System > Certificate**.

**Step 6** In the **Upload Certificate** area, click the file selection button. In the window for selecting files, select the obtained **.tar** certificate file packages and open them and click **Upload**. The system automatically imports the certificate.

**Step 7** After the certificate is imported, the system prompts you to synchronize the cluster configuration and restart the web service for the new certificate to take effect. After you complete these operations, click **OK**.

**Step 8** In the dialog box that is displayed, enter the password and click **OK** to automatically synchronize the cluster configuration and restart the web service.

**Step 9** After the cluster is restarted, enter the URL for accessing FusionInsight Manager in the address box of the browser and check whether the FusionInsight Manager web UI can be successfully displayed.

#### NOTE

The enterprise certificate has expired or security is hardened. After replacing the MRS certificate, replace the local certificate as well.

**Step 10** Log in to FusionInsight Manager and choose **Cluster > Overview > More > Restart**. In the displayed dialog box, enter the password of the current login user and click **OK**.

 NOTE

After the CA certificate is replaced, you need to restart the cluster offline to make the certificate take effect. Rolling restart is not supported.

**Step 11** In the displayed restart confirmation dialog box, click **OK**.

----End

## 7.10.2.2 Replacing an HA Certificate

HA certificates are used to encrypt the communication between active/standby processes and HA processes to ensure the communication security. This section describes how to replace the HA certificates on the active and standby management nodes on Manager to ensure product security. This feature applies to the following scenarios:

- After a cluster is installed for the first time, import an enterprise certificate.
- If an enterprise certificate has expired or requires security hardening, replace it with a new one.

You cannot replace the HA certificate for a cluster that does not have active and standby management nodes installed.

You can either apply for the certificate and key files from your enterprise certificate center or generate them yourself as a cluster user.

## Impact on the System

During the replacement process, Manager needs to be restarted, which will result in the system being inaccessible and unable to provide services.

## Prerequisites

- You have obtained the **root-ca.crt** HA root certificate file and the **root-ca.pem** key file to be replaced.
- You have prepared a password, such as **Userpwd@123**, for accessing the key file.

To avoid potential security risks, the password must meet the following complexity requirements:

- The password must contain at least eight characters.
- The password contains at least four types of the following: uppercase letters, lowercase letters, numbers, and special characters (`~`!?,;-'(){}[]/<>@#$$%^&*+|\|=`).
- When applying for a certificate from the certificate center, please provide the password to access the key file and request certificate files in formats such as CRT, CER, CERT, and PEM, as well as KEY and PEM format key files. The certificate you apply for needs to be able to issue certificates.

## Replacing an HA Certificate (MRS 3.x or Later)

**Step 1** Log in to the active management node as user **omm**.

**Step 2** Select a way to generate certificate and key files.

- If you choose to apply for the certificate and key files from your enterprise certificate center, save the certificate and key files to the **\$ {OMS\_RUN\_PATH}/workspace0/ha/local/cert** directory on the active and standby management nodes.

 **NOTE**

If the obtained certificate file is not in the **.crt** format or the key file is not in the **.pem** format, run either of the following commands accordingly to correct the format:

```
mv Certificate file name.Certificate file format root-ca.crt
```

```
mv Key file name.Key file format root-ca.pem
```

For example, run the following commands to name the certificate file **root-ca.crt** and the key file **root-ca.pem**:

```
mv server.cer root-ca.crt
```

```
mv server_key.key root-ca.pem
```

- If you choose to generate them yourself as a cluster user, run the following command to generate the **root-ca.crt** and **root-ca.pem** files in the **\$ {OMS\_RUN\_PATH}/workspace0/ha/local/cert** directory on the active management node:

```
sh ${OMS_RUN_PATH}/workspace/ha/module/hacom/script/gen-cert.sh --
root-ca --country=CN --state=state --city=city --company=company --
organize=organize --common-name=commonname --email=Cluster user
email address
```

 **NOTE**

The generated certificate files are valid for 10 years. When the system certificate files are about to expire, an alarm with the message "ALM-12055 Certificate File Is About to Expire" will be generated.

For example, run the following command:

```
sh ${OMS_RUN_PATH}/workspace/ha/module/hacom/script/gen-cert.sh --
root-ca --country=CN --state=guangdong --city=shenzhen --
company=huawei --organize=IT --common-name=HADOOP.COM --
email=abc@example.com
```

Enter the password as prompted and press **Enter**.

```
Enter pass phrase for /opt/huawei/Bigdata/om-server/OMS/workspace/ha/local/cert/root-ca.pem:
```

The command is successfully executed if the following information is displayed:

```
Generate root-ca pair success.
```

**Step 3** On the active management node, run the following command as user **omm** to copy **root-ca.crt** and **root-ca.pem** to the  **\${BIGDATA\_HOME}/om-server/om/security/certHA** directory:

```
cp -arp ${OMS_RUN_PATH}/workspace0/ha/local/cert/root-ca.* $
{BIGDATA_HOME}/om-server/om/security/certHA
```

**Step 4** Copy **root-ca.crt** and **root-ca.pem** generated on the active management node to the  **\${BIGDATA\_HOME}/om-server/om/security/certHA** directory on the standby management node as user **omm**.



```
scp ${OMS_RUN_PATH}/workspace0/ha/local/cert/root-ca.* omm@IP address
of the standby management node:${BIGDATA_HOME}/om-server/om/security/
certHA
```

- Step 5** Run the following command to generate an HA certificate and perform the automatic replacement:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/replacehaSSLCert.sh
```

Enter the password as prompted and press **Enter**.

```
Please input ha ssl cert password:
```

The HA certificate is replaced successfully if the following information is displayed:

```
[INFO] Succeed to replace ha ssl cert.
```

#### NOTE

If you need to update the HA password encryption suite, use the **-u** parameter.

- Step 6** Run the following command to restart OMS:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/restart-oms.sh
```

The following information is displayed:

```
start HA successfully.
```

- Step 7** Log in to the standby management node as user **omm** using the IP address of the standby management node, and repeat steps [Step 5](#) and [Step 6](#).

Run `sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh` to check whether **HAAllResOK** of the management node is **Normal** and whether FusionInsight Manager can be relogged in to. If they are, the operation is successful.

----End

## Replacing an HA Certificate (MRS 2.x or Earlier)

- Step 1** Log in to the active management node.

- Step 2** Run the following commands to switch the user:

```
sudo su - root
```

```
su - omm
```

- Step 3** Run the following commands to generate **root-ca.crt** and **root-ca.pem** in the `${OMS_RUN_PATH}/workspace0/ha/local/cert` directory on the active management node:

```
sh ${OMS_RUN_PATH}/workspace/ha/module/hacom/script/gen-cert.sh --
root-ca --country=country --state=state --city=city --company=company --
organize=organize --common-name=commonname --email=Administrator email
address --password=password
```

There can be security risks if a command contains the authentication password. You are advised to disable the command recording function (history) before running the command.

For example, run the following command: `sh ${OMS_RUN_PATH}/workspace/ha/module/hacom/script/gen-cert.sh --root-ca --country=CN --`

```
state=gd --city=sz --company=hw --organize=IT --common-
name=HADOOP.COM --email=abc@example.com --password=xxx
```

The command has been executed successfully if the following information is displayed:

```
Generate root-ca pair success.
```

- Step 4** On the active management node, run the following command as user **omm** to copy **root-ca.crt** and **root-ca.pem** to the `${BIGDATA_HOME}/om-0.0.1/security/certHA` directory:

```
cp -arp ${OMS_RUN_PATH}/workspace0/ha/local/cert/root-ca.* $
{BIGDATA_HOME}/om-0.0.1/security/certHA
```

- Step 5** Copy **root-ca.crt** and **root-ca.pem** generated on the active management node to the `${BIGDATA_HOME}/om-0.0.1/security/certHA` directory on the standby management node as user **omm**.

- Step 6** Run the following command to generate an HA certificate and perform the automatic replacement:

```
sh ${BIGDATA_HOME}/om-0.0.1/sbin/replacehaSSLCert.sh
```

Enter the password as prompted, and press **Enter**.

```
Please input ha ssl cert password:
```

The HA certificate is replaced successfully if the following information is displayed:

```
[INFO] Succeed to replace ha ssl cert.
```

- Step 7** Run the following command to restart OMS:

```
sh ${BIGDATA_HOME}/om-0.0.1/sbin/restart-oms.sh
```

The following information is displayed:

```
start HA successfully.
```

- Step 8** Log in to the standby management node and switch to user **omm**. Repeat step [Step 6](#) to step [Step 7](#).

Run the `sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh` command to check whether **HAAllResOK** of the management node is **Normal**. Access MRS Manager again. If MRS Manager can be accessed, the operation is successful.

----End

## 7.10.3 MRS Cluster Security Hardening

### 7.10.3.1 MRS Cluster Security Hardening Policies

#### Hardening Tomcat

Tomcat is hardened as follows based on open-source software during FusionInsight Manager software installation and use:

- The Tomcat version is upgraded to the official version.

- Permissions on the directories under applications are set to **500**, and the write permission on some directories is supported.
- The Tomcat installation package is automatically deleted after the system software is installed.
- The automatic deployment function is disabled for projects in application directories. Only the **web**, **cas**, and **client** projects are deployed.
- Some unused **http** methods are disabled, preventing attacks by using the **http** methods.
- The default shutdown port and command of the Tomcat server are changed to prevent hackers from shutting down the server and attacking servers and applications.
- To ensure security, the value of **maxHttpHeaderSize** is changed, which enables server administrators to control abnormal requests of clients.
- The Tomcat version description file is modified after Tomcat is installed.
- To prevent disclosure of Tomcat information, the Server attributes of Connector are modified so that attackers cannot obtain information about the server.
- Permissions on files and directories of Tomcat, such as the configuration files, executable files, log directories, and temporary folders, are under control.
- Session facade recycling is disabled to prevent request leakage.
- LegacyCookieProcessor is used as CookieProcessor to prevent the leakage of sensitive data in cookies.

## Hardening LDAP

LDAP is hardened as follows after a cluster is installed:

- In the LDAP configuration file, the password of the administrator account is encrypted using SHA. After the OpenLDAP is upgraded to 2.4.39 or later, data is automatically synchronized between the active and standby LDAP nodes using the SASL External mechanism, which prevents disclosure of the password.
- The LDAP service in the cluster supports the SSLv3 protocol by default, which can be used safely. When the OpenLDAP is upgraded to 2.4.39 or later, the LDAP automatically uses TLS1.0 or later to prevent unknown security risks.

## Hardening JDK

- If the client process uses the AES256 encryption algorithm, JDK security hardening is required. The operations are as follows:

Obtain the Java Cryptography Extension (JCE) package whose version matches that of JDK. The JCE package contains **local\_policy.jar** and **US\_export\_policy.jar**. Copy the JAR files to the following directory and replace the files in the directory.

- Linux: *JDK installation directory*/jre/lib/security
- Windows: *JDK installation directory*\jre\lib\security

### NOTE

Access the Open JDK open-source community to obtain the JCE file.

- If the client process uses the SM4 encryption algorithm, the JAR package needs to be updated.

Obtain **SMS4JA.jar** in the *client installation directory/JDK/jdk/jre/lib/ext/* directory, and copy the JAR package to the following directory:

- Linux: *JDK installation directory/jre/lib/ext/*
- Windows: *JDK installation directory\jre\lib\ext\*

### 7.10.3.2 Configuring Hadoop Data Encryption During Transmission

#### Configuring Security Channel Encryption

The channels between components are not encrypted by default. You can set the following parameters to configure security channel encryption.

To modify parameters, log in to FusionInsight Manager, choose **Cluster > Services > Target Service Name**, and click **Configurations** then **All Configurations**. Enter a parameter name in the search box.

#### NOTE

- Restart corresponding services for the modification to take effect after you modify configuration parameters.
- This topic is available for MRS 3.x or later.

Table 7-78 Parameter description

| Service | Parameter                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Default Value                                                                                                                |
|---------|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| HBase   | hbase.rpc.protection      | <p>Indicates whether the HBase channels, including the remote procedure call (RPC) channels for HBase clients to access the HBase server and the RPC channels between the HMaster and RegionServer, are encrypted. If this parameter is set to <b>privacy</b>, the channels are encrypted and the authentication, integrity, and privacy functions are enabled. If this parameter is set to <b>integrity</b>, the channels are not encrypted and only the authentication and integrity functions are enabled. If this parameter is set to <b>authentication</b>, the channels are not encrypted, only packets are authenticated, and integrity and privacy are not required.</p> <p><b>NOTE</b><br/>The privacy mode encrypts transmitted content, including sensitive information such as user tokens, to ensure the security of the transmitted content. However, this mode has great impact on performance. Compared with the other two modes, this mode reduces read/write performance by about 60%. Modify the configuration based on the enterprise security requirements. The configuration items on the client and server must be the same.</p> | <ul style="list-style-type: none"><li>• Security mode: <b>privacy</b></li><li>• Normal mode: <b>authentication</b></li></ul> |
| HDFS    | dfs.encrypt.data.transfer | <p>Indicates whether the HDFS data transfer channels and the channels for clients to access HDFS are encrypted. The HDFS data transfer channels include the data transfer channels between DataNodes and the Data Transfer (DT) channels for clients to access DataNodes. The value <b>true</b> indicates that the channels are encrypted. The channels are not encrypted by default.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | false                                                                                                                        |

| Service | Parameter                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Default Value                                                                                                                   |
|---------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| HDFS    | dfs.encrypt.data.transfer.algorithm | <p>Indicates the encryption algorithm of the HDFS data transfer channels and the channels for clients to access HDFS. This parameter is available only when <b>dfs.encrypt.data.transfer</b> is set to <b>true</b>.</p> <p>The default value is <b>3des</b>, indicating that 3DES algorithm is used to encrypt data. The value can also be set to <b>rc4</b>. However, to avoid security risks, you are not advised to set the parameter to this value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 3des                                                                                                                            |
| HDFS    | hadoop.rpc.protection               | <p>Indicates whether the RPC channels of each module in Hadoop are encrypted. The channels include:</p> <ul style="list-style-type: none"> <li>• RPC channels for clients to access HDFS</li> <li>• RPC channels between HDFS modules, for example, between DataNode and NameNode</li> <li>• RPC channels for clients to access YARN</li> <li>• RPC channels between NodeManager and ResourceManager</li> <li>• RPC channels for Spark to access YARN and HDFS</li> <li>• RPC channels for MapReduce to access YARN and HDFS</li> <li>• RPC channels for HBase to access HDFS</li> </ul> <p>The default value is <b>privacy</b>, indicating encrypted transmission. The value <b>authentication</b> indicates that transmission is not encrypted.</p> <p><b>NOTE</b><br/>You can set this parameter on the HDFS component configuration page. The parameter setting is valid globally, that is, the setting of whether the RPC channel is encrypted takes effect on all modules in Hadoop.</p> | <ul style="list-style-type: none"> <li>• Security mode: <b>privacy</b></li> <li>• Normal mode: <b>authentication</b></li> </ul> |

## Setting the Maximum Number of Concurrent Web Connections

To ensure web server reliability, new connections are rejected when the number of user connections reaches a specific threshold. This prevents DDOS attacks and

service unavailability caused by too many users accessing the web server at the same time.

To modify parameters, log in to FusionInsight Manager, choose **Cluster > Services > Target Service Name**, and click **Configurations** then **All Configurations**. Enter a parameter name in the search box.

**Table 7-79** Parameter description

| Service    | Parameter                      | Description                                                                   | Default Value |
|------------|--------------------------------|-------------------------------------------------------------------------------|---------------|
| HD FS/Yarn | hadoop.http.server.MaxRequests | Specifies the maximum number of concurrent web connections of each component. | 2000          |
| Spark2x    | spark.connection.maxRequest    | Specifies the maximum number of request connections of JobHistory.            | 5000          |

### 7.10.3.3 Configuring Kafka Data Encryption During Transmission

#### Scenario

Data between the Kafka client and the broker is transmitted in plain text. The Kafka client may be deployed in an untrusted network, exposing the transmitting data to leakage and tampering risks.

This topic is available for MRS 3.x or later.

#### Configuring Kafka Data Encryption During Transmission

The channel between components is not encrypted by default. You can set the following parameters to enable security channel encryption.

To modify parameters, log in to FusionInsight Manager, choose **Cluster > Services > Kafka**, click **Configurations** then **All Configurations**, and enter a parameter name in the search box.

#### NOTE

After the configuration, restart the corresponding service for the settings to take effect.

**Table 7-80** describes the parameters related to transmission encryption on the Kafka server.

**Table 7-80** Parameters relevant to Kafka data encryption during transmission

| Parameter                      | Description                                                                                                                                                                                 | Default Value  |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| ssl.mode.enable                | Indicates whether to enable the Secure Sockets Layer (SSL) protocol. If this parameter is set to <b>true</b> , services relevant to the SSL protocol are started during the broker startup. | false          |
| security.inter.broker.protocol | Indicates communication protocol between brokers. The communication protocol can be PLAINTEXT, SSL, SASL_PLAINTEXT, or SASL_SSL.                                                            | SASL_PLAINTEXT |

The SSL protocol can be configured for the server or client to encrypt transmission and communication only after **ssl.mode.enable** is set to **true** and broker enables the **SSL** and **SASL\_SSL** protocols.

### 7.10.3.4 Configuring HDFS Data Encryption During Transmission

This section describes how to configure encryption for HDFS security channels.

This topic is available for MRS 3.x or later.

## Configuring HDFS Security Channel Encryption

The channel between components is not encrypted by default. You can set parameters to enable security channel encryption.

To modify parameters, log in to FusionInsight Manager, choose **Cluster > Services > HDFS**, and click **Configurations** then **All Configurations**. Enter a parameter name in the search box.

#### NOTE

After the configuration, restart the corresponding service for the settings to take effect.



**Table 7-81** Parameters

| Configuration Item    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Default Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hadoop.rpc.protection | <p><b>NOTICE</b></p> <ul style="list-style-type: none"> <li>The setting takes effect only after the service is restarted. Rolling restart is not supported.</li> <li>After the setting, you need to download the client configuration file again. Otherwise, HDFS cannot provide the read and write services.</li> <li>After the setting, you need to restart the executor. Otherwise, the job management and file management functions on the console become unavailable.</li> </ul> <p>Indicates whether the RPC channels of each module in Hadoop are encrypted. The channels include:</p> <ul style="list-style-type: none"> <li>RPC channels for clients to access HDFS</li> <li>RPC channels between modules in HDFS, for example, between DataNode and NameNode</li> <li>RPC channels for clients to access Yarn</li> <li>RPC channels between NodeManager and ResourceManager</li> <li>RPC channels for Spark to access Yarn and HDFS</li> <li>RPC channels for MapReduce to access Yarn and HDFS</li> <li>RPC channels for HBase to access HDFS</li> </ul> <p><b>NOTE</b><br/>The setting takes effect globally, that is, the encryption attribute of the RPC channel of each module in the Hadoop takes effect.</p> | <ul style="list-style-type: none"> <li>Security mode: privacy</li> <li>Normal mode: authentication</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li><b>authentication:</b> indicates that only authentication is required.</li> <li><b>integrity:</b> indicates that authentication and consistency check need to be performed.</li> <li><b>privacy:</b> indicates that authentication, consistency check, and encryption need to be performed.</li> </ul> |

| Configuration Item                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Default Value         |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| dfs.encrypt.data.transf<br>er               | <p>Indicates whether the HDFS data transfer channels and the channels for clients to access HDFS are encrypted. The HDFS data transfer channels include the data transfer channels between DataNodes and the Data Transfer (DT) channels for clients to access DataNodes. The value <b>true</b> indicates that the channels are encrypted. The channels are not encrypted by default.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• This parameter is available only when <b>hadoop.rpc.protection</b> is set to <b>privacy</b>.</li> <li>• If a large amount of service data is transmitted, enabling encryption by default severely affects system performance.</li> <li>• If data transmission encryption is configured for one cluster in the trusted cluster, the same data transmission encryption must be configured for the peer cluster.</li> </ul> | false                 |
| dfs.encrypt.data.transf<br>er.algorithm     | <p>Indicates the algorithm to encrypt the HDFS data transfer channels and the channels for clients to access HDFS. This parameter is available only when <b>dfs.encrypt.data.transfer</b> is set to <b>true</b>.</p> <p><b>NOTE</b></p> <p>The default value is <b>3des</b>, indicating that 3DES algorithm is used to encrypt data. The value can also be set to <b>rc4</b>. However, to avoid security risks, you are not advised to set the parameter to this value.</p>                                                                                                                                                                                                                                                                                                                                                                                                    | 3des                  |
| dfs.encrypt.data.transf<br>er.cipher.suites | <p>This parameter can be left empty or set to <b>AES/CTR/NoPadding</b> to specify the cipher suite for data encryption. If this parameter is not specified, the encryption algorithm specified by <b>dfs.encrypt.data.transfer.algorithm</b> is used for data encryption. The default value is <b>AES/CTR/NoPadding</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | AES/CTR/<br>NoPadding |

### 7.10.3.5 Configuring Spark Data Encryption During Transmission

#### Scenario

This section describes how to configure encryption for Spark security channels.

This topic is available for MRS 3.x or later.

#### Configuring Spark Data Encryption During Transmission

To modify parameters, log in to Manager, choose **Cluster > Services > Spark**, click **Configurations** then **All Configurations**, and enter a parameter name in the search box.

#### NOTE

After the configuration, restart the corresponding service for the settings to take effect.

**Table 7-82** Parameters

| Parameter                                | Description                                                                                         | Default Value                                           |
|------------------------------------------|-----------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| spark.authenticate                       | Whether to enable Spark internal security authentication                                            | Security mode: <b>true</b><br>Normal mode: <b>false</b> |
| spark.authenticate.enableSaslEncryption  | Whether to enable encrypted communication based on Simple Authentication and Security Layer (SASL). | Security mode: <b>true</b><br>Normal mode: <b>false</b> |
| spark.network.crypto.enabled             | Whether to enable RPC encryption based on Advanced Encryption Standard (AES)                        | Security mode: <b>true</b><br>Normal mode: <b>false</b> |
| spark.network.sasl.serverAlwaysEncrypt   | Whether to disable unencrypted connections for ports with SASL authentication enabled               | false                                                   |
| spark.network.crypto.keyLength           | Length of the encryption key to be generated                                                        | 256                                                     |
| spark.network.crypto.keyFactoryAlgorithm | Algorithm used to generate the encryption key                                                       | PBKDF2WithHmacSHA1                                      |

| Parameter                            | Description                                                                           | Default Value                                           |
|--------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------|
| spark.io.encryption.enabled          | Whether to enable local disk I/O encryption.                                          | Security mode: <b>true</b><br>Normal mode: <b>false</b> |
| spark.io.encryption.keygen.algorithm | Algorithm used to generate the I/O encryption key                                     | HmacSHA256                                              |
| spark.io.encryption.keySizeBits      | Size of an I/O encryption key, in bits                                                | 256                                                     |
| spark.ssl.ui.enabled                 | Whether to enable Secure Sockets Layer (SSL) authentication for the web UI connection | Security mode: <b>true</b><br>Normal mode: <b>false</b> |

### 7.10.3.6 Configuring ZooKeeper Data Encryption During Transmission

#### Scenario

By default, SSL channel encryption transmission is disabled between the ZooKeeper client and server and between instances on the server. This section describes how to enable the ZooKeeper channel encryption transmission.

#### NOTE

This function is available only for MRS clusters of version 3.1.2 or later.

#### Impact on the System

- When SSL channel encryption transmission is enabled on the ZooKeeper server, the performance deteriorates.
- When SSL channel encryption transmission is enabled on the ZooKeeper server, ZooKeeper and dependent upper-layer components need to be restarted. During the restart, services are unavailable.
- To enable SSL channel encryption transmission on the ZooKeeper server, you need to download the client again.
- If SSL channel encryption transmission is enabled for ZooKeeper, rolling restart is not supported.

### Configuring ZooKeeper Data Encryption During Transmission

**Step 1** Log in to FusionInsight Manager, click **Cluster** and choose **Services > ZooKeeper**. On the displayed page, click **Configurations** and click **All Configurations**.

**Step 2** Enter the parameter name in the search box, and change the value as follows:

**Table 7-83** Security configuration item

| Parameter   | Description                                     | Default Value | New Value |
|-------------|-------------------------------------------------|---------------|-----------|
| ssl.enabled | Whether to enable SSL communication encryption. | false         | true      |

**Step 3** After the modification is complete, click **Save** and then click **OK**.

**Step 4** Click **Dashboard**. In the upper right corner of the displayed page, choose **More > Restart Service**, enter the password, and confirm the operation impact.

You can select **Restart upper-layer services**. During the restart of all affected components, services will be unavailable. Exercise caution when performing this operation.

**Step 5** Click **OK** and wait until the services are restarted successfully.

**Step 6** Choose **Cluster > Active/Standby Cluster DR** to check whether active/standby DR is configured for the current cluster.

- If yes, go to **Step 7**.
- If no, no further action is required.

**Step 7** The **ssl.enabled** configuration of the ZooKeeper service in the active cluster must be the same as that in the DR cluster. Modify the **ssl.enabled** parameter in the cluster where no operation is performed by referring to the preceding steps.

**Step 8** Log in to the active OMS node in the active cluster as user **root** and run the following commands to restart the DR management process:

```
su - omm
```

```
`${BIGDATA_HOME}/om-server/om/share/om/disaster/sbin/restart-disaster.sh
```

If the following information is displayed, the operation is successful:

```
...
disaster start with process id : 23256
End into restart-disaster.sh
```

**Step 9** Log in to the active OMS node in the DR cluster as user **root** and run the following commands to restart the DR management process:

```
su - omm
```

```
`${BIGDATA_HOME}/om-server/om/share/om/disaster/sbin/restart-disaster.sh
```

```
----End
```

### 7.10.3.7 Encrypting Data Transmission Between the Controller and Agent

#### Scenario

After a cluster is installed, Controller and Agent need to communicate with each other. The Kerberos authentication is used during the communication. By default,

the communication is not encrypted during the communication for the sake of cluster performance. Users who have demanding security requirements can use the method described in this section for encryption.

#### NOTE

This topic is available for MRS 3.x or later.

## Impact on the System

- Controller and all Agents automatically restart, which interrupts FusionInsight Manager.
- The performance of management nodes deteriorates in large clusters. You are advised to enable the encryption function for clusters with a maximum of 200 nodes.

## Prerequisites

You have obtained the IP addresses of the active and standby management nodes.

## Encrypting Data Transmission Between the Controller and Agent

**Step 1** Log in to the active management node as user **omm**.

**Step 2** Run the following command to disable logout upon timeout:

```
TMOUT=0
```

#### NOTE

After the operations in this section are complete, run the **TMOUT=Timeout interval** command to restore the timeout interval in a timely manner. For example, **TMOUT=600** indicates that a user is logged out if the user does not perform any operation within 600 seconds.

**Step 3** Run the following command to go to the related directory:

```
cd ${CONTROLLER_HOME}/sbin
```

**Step 4** Run the following command to enable communication encryption:

```
./enableRPCencrypt.sh -t
```

Run the **sh \${BIGDATA\_HOME}/om-server/om/sbin/status-oms.sh** command to check whether **ResHASStatus** of the active management node Controller is **Normal** and whether you can log in to FusionInsight Manager again. If yes, the enablement is successful.

**Step 5** Run the following command to disable communication encryption when necessary:

```
./enableRPCencrypt.sh -f
```

Run the **sh \${BIGDATA\_HOME}/om-server/om/sbin/status-oms.sh** command to check whether **ResHASStatus** of the active management node Controller is **Normal** and whether you can log in to FusionInsight Manager again. If yes, the enablement is successful.

----End

## 7.10.3.8 Configuring a Trusted IP Address to Access LDAP

### Scenario

By default, the LDAP service deployed in the OMS and cluster can be accessed by any IP address. To enable the LDAP service to be accessed by only trusted IP addresses, you can configure the INPUT policy in the iptables filtering list.

#### NOTE

This topic is available for MRS 3.x or later.

### Impact on the System

After the configuration, the LDAP service cannot be accessed by IP addresses that are not configured. Before the expansion, the added IP addresses need to be configured as trusted IP addresses.

### Prerequisites

- You have collected the management plane IP addresses and service plane IP addresses of all nodes in the cluster and all floating IP addresses.
- You have obtained the **root** user account for all nodes in the cluster.

## Configuring a Trusted IP Address to Access LDAP

### Configuring trusted IP addresses for the LDAP service on the OMS

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **System > OMS** and choose **oldap > Modify Configuration** to view the OMS LDAP port number, that is, the value of **LDAP Listening Port**, which is 21750 by default.
- Step 3** Log in to the active management node as user **root** using the IP address of the active management node.
- Step 4** Run the following command to check the INPUT policy in the iptables filtering list:

#### **iptables -L**

For example, if no rule is configured, the INPUT policy is displayed as follows:

```
Chain INPUT (policy ACCEPT)
target prot opt source destination
```

- Step 5** Run the following command to trust all IP addresses used by the cluster. Each IP address needs to be added once.

```
iptables -A INPUT -s Trusted IP address -p tcp --dport Port number -j ACCEPT
```

For example, to configure **10.0.0.1** as a trusted IP address and enable it to access port **21750**, you need to run the following command:

```
iptables -A INPUT -s 10.0.0.1 -p tcp --dport 21750 -j ACCEPT
```

- Step 6** Run the following command to untrust all IP addresses. Trusted IP addresses are not affected by this rule.

**iptables -A INPUT -p tcp --dport *Port number* -j DROP**

For example, to disable all IP addresses to access port **21750**, run the following command:

**iptables -A INPUT -p tcp --dport 21750 -j DROP**

- Step 7** Run the following command to view the modified INPUT policy in the iptables filtering list:

**iptables -L**

For example, after a trusted IP address is configured, the INPUT policy is displayed as follows:

```
Chain INPUT (policy ACCEPT)
target prot opt source destination
ACCEPT tcp -- 10.0.0.1 anywhere tcp dpt:21750
DROP tcp -- anywhere anywhere tcp dpt:21750
```

- Step 8** Run the following command to view the rules and rule numbers in the iptables filtering list:

**iptables -L -n --line-number**

```
Chain INPUT (policy ACCEPT)
num target prot opt source destination
1 DROP tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:21750
```

- Step 9** Run the following command to delete the desired rule from the iptables filtering list based on site requirement:

**iptables -D INPUT *Number of the rule to be deleted***

For example, to delete rule 1, run the following command:

**iptables -D INPUT 1**

- Step 10** Log in to the standby management node as user **root** using the standby IP address. Repeat [Step 4](#) to [Step 9](#).

### Configuring trusted IP addresses for the LDAP service in the cluster

- Step 11** Log in to FusionInsight Manager.

- Step 12** Choose **Cluster > Services > LdapServer**. Click **Instance** and view the LDAP nodes.

- Step 13** Go to the **Configurations** page, and view the LDAP port number of the cluster, that is, the value of **LDAP\_SERVER\_PORT**, which is 21780 by default.

- Step 14** Log in to the LDAP node as user **root** using the LDAP service IP address.

- Step 15** Run the following command to view the INPUT policy in the iptables filtering list:

**iptables -L**

For example, if no rule is configured, the INPUT policy is displayed as follows:

```
Chain INPUT (policy ACCEPT)
target prot opt source destination
```

- Step 16** Run the following command to configure all IP addresses used by the cluster as trusted IP addresses. Each IP address needs to be added independently.

**iptables -A INPUT -s *Trusted IP address* -p tcp --dport *Port number* -j ACCEPT**



For example, to configure **10.0.0.1** as a trusted IP address and enable it to access port **21780**, you need to run the following command:

```
iptables -A INPUT -s 10.0.0.1 -p tcp --dport 21780 -j ACCEPT
```

- Step 17** Run the following command to untrust all IP addresses. Trusted IP addresses are not affected by this rule.

```
iptables -A INPUT -p tcp --dport Port number -j DROP
```

For example, to disable all IP addresses to access port **21780**, run the following command:

```
iptables -A INPUT -p tcp --dport 21780 -j DROP
```

- Step 18** Run the following command to view the modified INPUT policy in the iptables filtering list:

```
iptables -L
```

For example, after a trusted IP address is configured, the INPUT policy is displayed as follows:

```
Chain INPUT (policy ACCEPT)
target prot opt source destination tcp dpt:21780
ACCEPT tcp -- 10.0.0.1 anywhere tcp dpt:21780
DROP tcp -- anywhere anywhere tcp dpt:21780
```

- Step 19** Run the following command to view the rules and rule numbers in the iptables filtering list:

```
iptables -L -n --line-number
```

```
Chain INPUT (policy ACCEPT)
num target prot opt source destination tcp dpt:21780
1 DROP tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:21780
```

- Step 20** Run the following command to delete the desired rule from the iptables filtering list based on site requirement:

```
iptables -D INPUT Number of the rule to be deleted
```

For example, to delete rule 1, run the following command:

```
iptables -D INPUT 1
```

- Step 21** Log in to the LDAP node as user **root** using the IP address of another LDAP service, and repeat [Step 15](#) to [Step 20](#).

----End

### 7.10.3.9 HFile and WAL Encryption

HFile and Write ahead log (WAL) in HBase are not encrypted by default. To encrypt them, perform the operations provided in this topic.

**NOTICE**

- Setting the HFile and WAL encryption mode to SMS4 or AES has a great impact on the system and will cause data loss in case of any misoperation. You are not advised to perform this operation.
- Batch data import using BulkLoad does not support data encryption.
- This topic is available for MRS 3.x and later versions only.

## HFile and WAL Encryption

**Step 1** On any HBase node, run the following commands to create a key file as user **omm**:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-
HBase-2.2.3/hbase/bin/hbase-encrypt.sh <path>/hbase.jks <type> <length>
<alias>
```

- *<path>/hbase.jks* indicates the path for storing the generated JKS file.
- *<type>* indicates the encryption type, which can be SMS4 and AES.
- *<length>* indicates the key length. SMS4 supports 16-bit and AES supports 128-bit.
- *<alias>* indicate the alias of the key file. When you create the key file for the first time, retain the default value **omm**.

For example, to generate an SMS4 encryption key, run the following command:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-
HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks SMS4 16
omm
```

To generate an AES encryption key, run the following command:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-
HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks AES 128
omm
```

**NOTE**

- To ensure operations can be successfully performed, the *<path>/hbase.jks* directory needs to be created in advance, and the cluster operation user must have the **rw** permission of this directory.
- After running the command, enter the same *<password>* for four times. The password is the same as the one encrypted in [Step 3](#).

**Step 2** Distribute the generated key files to the same directory on all nodes in the cluster and assign read and write permission to user **omm**.

**NOTE**

- Administrators need to select a safe procedure to distribute keys based on the enterprise security requirements.
- If the key files of some nodes are lost, repeat the step to copy the key files from other nodes.

**Step 3** On FusionInsight Manager, choose **Cluster > Services > HBase > Configurations**. Search for and configure the following parameters:

- **hbase.crypto.keyprovider.parameters.encryptedtext**: Set this parameter to a ciphertext password in the format of `<encrypted_password>`.

`<encrypted_password>` indicates the encrypted password generated during the key file creation. The parameter value is displayed in ciphertext. Run the following command as user **omm** to obtain the related encrypted password on the nodes where HBase service is installed:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh
```

 **NOTE**

Enter the `<password>`, which is the same password you entered in [Step 1](#).

- **hbase.crypto.keyprovider.parameters.uri**: Set this parameter to the key path and name in the `jceks:// <key_Path_Name>` format.  
`<key_Path_Name>` indicates the path of the key file. For example, if the path of the key file is `/home/hbase/conf/hbase.jks`, set this parameter to `jceks:///home/hbase/conf/hbase.jks`.
- **hbase.crypto.key.algorithm**: If this parameter is set to **SMS4** or **AES**, HFile content is encrypted using SMS4 or AES.
- **hbase.crypto.wal.algorithm**: If this parameter is set to **SMS4** or **AES**, WAL content is encrypted in SMS4 or AES mode.
- **hbase.regionserver.wal.encryption**: Set this parameter to **true**.

**Step 4** Click **Save**. Click **Dashboard**. In the upper right corner of the page, choose **More > Restart Service**, enter the password of the current user, and click **OK** to apply the changes.

**Step 5** Create an HBase table through CLI or code and configure the encryption mode to enable encryption. `<type>` indicates the encryption type, and **d** indicates the column family.

- When you create an HBase table through CLI, set the encryption mode to SMS4 or AES for the column family.

```
create '<table name>', {NAME => 'd', ENCRYPTION => '<type>'}
```

- When you create an HBase table using code, set the encryption mode to SMS4 or AES by adding the following information to the code:

```
public void testCreateTable()
{
 String tableName = "user";
 Configuration conf = getConfiguration();
 HTableDescriptor htd = new HTableDescriptor(TableName.valueOf(tableName));

 HColumnDescriptor hcd = new HColumnDescriptor("d");
 //Set the encryption mode to SMS4 or AES.
 hcd.setEncryptionType("<type>");
 htd.addFamily(hcd);

 HBaseAdmin admin = null;
 try
 {
 admin = new HBaseAdmin(conf);

 if(!admin.tableExists(tableName))
 {
 admin.createTable(htd);
 }
 }
}
```



5. You can perform [Step 6](#) to check whether the encryption configuration is successful.

----End

## Verifying the Encryption Configuration

### NOTE

This operation can be performed only when test data can be written to an empty table.

- Step 1** Log in to the node where the client is installed as the client installation user. Switch to the client installation directory, for example, **/opt/client**.

```
cd /opt/client
```

- Step 2** Run the following command to set environment variables:

```
source bigdata_env
```

- Step 3** Run the following command to authenticate the current user if Kerberos authentication has been enabled for the current security cluster. The current user must have the permission to read and write HBase tables and the HDFS operation permission.

```
kinit Component service user
```

If Kerberos authentication is disabled for the cluster, set the Hadoop username.

```
export HADOOP_USER_NAME=hbase
```

- Step 4** Run the following command to log in to the HBase client:

```
hbase shell
```

Run the following command to insert a new data record and flush the table to generate an HFile:

```
put '<table_name>', 'fd2', 'dc1', 'value22222222222222222222222222222222'
```

```
flush '<table_name>'
```

### NOTE

- *<table\_name>* indicates the table configured with SMS4 or AES encryption. For details about how to configure SMS4 or AES encryption, go to [Step 5](#).
- *d* indicates the column family configured with SMS4 or AES encryption. For details about how to configure SMS4 or AES encryption, go to [Step 5](#).

- Step 5** Press **Ctrl+C** to exit the HBase client.

- Step 6** Run the following command to view the directory where the HFile file generated in [Step 4](#) is stored:

```
hdfs dfs -ls
```

The file directory format is **/hbase/data/<namespace\_name>/<table\_name>/<region\_name>/<columnfamily\_name>/<HFile\_name>**.

 NOTE

If `<namespace_name>` is not specified during HBase table creation, **default** is used by default.

Example:

```
/hbase/data/default/create_table/dd61b81b1ba1aad6513b9bdcd8f871c/d/
aa6fe387b27443afaba40f5b584c1fa7
```

**Step 7** Run the following command to view the HFile content:

```
hbase hfile -f <HFile path> -p
```

 NOTE

`<HFile path>` indicates the directory where the HFile file is located.

The error message "com.huawei.hadoop.hbase.io.crypto.CryptoRuntimeException" will be displayed in the command output. However, the **HBase shell** can still read the table data, indicating that the encryption configuration is successful.

----End

## Modifying a Key File

During the **HFile and WAL Encryption** operation, the related key file must be generated and its password must be set to ensure system security. After a period of running, you can replace the key file with a new one to encrypt HFile and WAL.

---

**NOTICE**

Modifying a key file has a great impact on the system and will cause data loss in case of any misoperation. You are not advised to perform this operation.

---

**Step 1** Run the following command to generate a new key file as user **omm**:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-
HBase-2.2.3/hbase/bin/hbase-encrypt.sh <path>/hbase.jks <type> <length>
<alias-new>
```

- `<path>/hbase.jks`: indicates the path for storing the generated **hbase.jks** file. The path and file name must be consistent with those of the key file generated in **HFile and WAL Encryption**.
- `<alias-new>`: indicates the alias of the key file. The alias must be different with that of the old key file.
- `<type>`: indicates the encryption type, which can be SMS4 or AES.
- `<length>` indicates the key length. SMS4 supports 16-bit and AES supports 128-bit.

For example, to generate an SMS4 encryption key, run the following command:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-
HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks SMS4 16
omm_new
```

To generate an AES encryption key, run the following command:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-
HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks AES 128
omm_new
```

 NOTE

- To ensure operations can be successfully performed, the `<path>/hbase.jks` directory needs to be created in advance, and the cluster operation user must have the `rw` permission of this directory.
- After running the command, you need to enter the same `<password>` for three times. This password is the password of the key file. You can use the password of the old file without any security risk.

**Step 2** Distribute the generated key files to the same directory on all nodes in the cluster and assign read and write permission to user `omm`.

 NOTE

Administrators need to select a safe procedure to distribute keys based on the enterprise security requirements.

**Step 3** Log in to FusionInsight Manager, choose **Cluster > Services > HBase > Configurations**, search for `hadoop.config.expandor` in the search box, and add the following custom parameters:

- The name of the new custom parameter is `hbase.crypto.master.key.name`, and the value is `omm_new`.
- The name of the new custom parameter is `hbase.crypto.master.alternate.key.name`, and the value is `omm`

**Figure 7-36** Adding a custom parameter

| Parameter                   | Value                                                                                                                                                                                                             |       |       |                            |         |                             |     |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|-------|----------------------------|---------|-----------------------------|-----|
| hadoop.config.expandor      | <table border="1"><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td>hbase.crypto.master.key.na</td><td>omm_new</td></tr><tr><td>hbase.crypto.master.alterna</td><td>omm</td></tr></tbody></table> | Name  | Value | hbase.crypto.master.key.na | omm_new | hbase.crypto.master.alterna | omm |
|                             | Name                                                                                                                                                                                                              | Value |       |                            |         |                             |     |
| hbase.crypto.master.key.na  | omm_new                                                                                                                                                                                                           |       |       |                            |         |                             |     |
| hbase.crypto.master.alterna | omm                                                                                                                                                                                                               |       |       |                            |         |                             |     |

**Step 4** Click **Save**. Click **Dashboard**. In the upper right corner of the page, choose **More > Restart Service**, enter the password of the current user, and click **OK** to apply the changes.

**Step 5** Log in to the Hbase client by referring to [Step 1~Step 4 in Verifying the Encryption Configuration](#) and run the `major compact` command to generate an HFile file based on the new encryption algorithm.

```
major_compact '<table_name>'
```

**Step 6** On FusionInsight Manager, choose **Cluster > Services > HBase** and click the hyperlink on the right of **HMaster Web UI**. In the **Region Servers** tab, click **Compactions** to view the `major compact` progress.

Region Servers

Base Stats Memory Requests Storefiles **Compactions** Replications

| ServerName    | Num. Compacting Cells | Num. Compacted Cells | Remaining Cells | Compaction Progress |
|---------------|-----------------------|----------------------|-----------------|---------------------|
| 1659665978456 | 3                     | 3                    | 0               | 100.00%             |
| 1659665978352 | 0                     | 0                    | 0               |                     |
| 1659665980089 | 2725                  | 2725                 | 0               | 100.00%             |
| 1659665981123 | 415                   | 415                  | 0               | 100.00%             |
| 1659665979991 | 29                    | 29                   | 0               | 100.00%             |
| 1659665979920 | 0                     | 0                    | 0               |                     |

**Step 7** When all items in **Compaction Progress** reach **100%** and those in **Remaining KVs** are **0**, run the following command as user **omm** to destroy the old key file:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh <path>/hbase.jks <alias-old>
```

- *<path>/hbase.jks*: indicates the path for storing the generated **hbase.jks** file. The path and file name must be consistent with those of the key file generated in **HFile and WAL Encryption**.
- *<alias-old>*: indicates the alias of the old key file to be deleted.

For example:

```
sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks omm
```

#### NOTE

To ensure operations can be successfully performed, the *<path>/hbase.jks* directory needs to be created in advance, and the cluster operation user must have the **rw** permission of this directory.

**Step 8** Repeat **Step 2** and distribute the updated key files again.

**Step 9** Delete the HBase self-defined configuration item **hbase.crypto.master.alternate.key.name** added in **Step 3** from FusionInsight Manager.

**Step 10** Repeat **Step 4** for the configuration take effect.

----End

### 7.10.3.10 Configuring the IP Address Whitelist for Modifying Data in an HBase Read-Only Cluster

If the Replication function is enabled for HBase in an MRS 3.x and later cluster, a protection mechanism for data modification is added on the standby HBase cluster to ensure data consistency between the active and standby clusters. Upon receiving an RPC request for data modification, the standby HBase cluster checks the permission of the user who sends the request (only HBase manage users have the modification permission). Then it checks the validity of the source IP address of the request. Only modification requests from IP addresses in the white list are accepted. The IP address white list is configured by the **hbase.replication.allowedIPs** item.

Log in to FusionInsight Manager and choose **Cluster > Services > HBase**. Click **Configurations** and search for the parameter name in **Table 7-84**.



**Table 7-84** Parameter description

| Parameter                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Default Value |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| hbase.replication.allowedIPs | <p>Only replication requests from specified IP addresses are allowed. Only regular expressions separated by commas (,) are supported. Each pattern can be any of the following:</p> <ul style="list-style-type: none"><li>• Regex pattern<br/>Example: 10.18.40.*, 10.18.*, 10.18.40.11</li><li>• Range pattern (Range can be specified only in the last octet)<br/>Example: 10.18.40.[10-20]</li></ul> <p>If this item is empty (default value), the white list contains only the IP address of the RegionServer of the cluster, indicating that only modification requests from the RegionServer of the standby HBase cluster are accepted.</p> | N/A           |

### 7.10.3.11 Configuring LDAP Output Audit Logs

Users can set the audit log output level of the LDAP service and output audit logs in a specified directory, for example, `/var/log/messages`. The logs output can be used to check user activities and operation commands.

#### NOTE

- Enabling LDAP audit log output can generate a large number of logs, impacting cluster performance. Use this feature with caution.
- This topic is available for MRS 3.x or later.

### Configuring the LDAP Firewall Policy

In the cluster adopting the dual-plane networking, the LDAP is deployed on the service plane. To ensure the LDAP data security, you are advised to configure the firewall policy in the cluster to disable relevant LDAP ports.

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster > Services > LdapServer** and click **Configurations**.
- Step 3** Check the value of **LDAP\_SERVER\_PORT**, which is the service port of LdapServer.
- Step 4** To ensure data security, configure the firewall policy for the whole cluster to disable the LdapServer port based on the customer's firewall environment.

----End

## Enabling the LDAP Audit Log Output

**Step 1** Log in to any LdapServer node.

**Step 2** Run the following command to edit the `slapd.conf.consumer` file, and set the value of `loglevel` to `256` (you can run the `man slapd.conf` command on the OS to view the log level definition).

```
cd ${BIGDATA_HOME}/FusionInsight_BASE_8.1.0.1/install/FusionInsight-
Ldapserver-2.7.0/ldapserver/local/template
```

```
vi slapd.conf.consumer
```

```
...
pidfile [PID_FILE_SLAPD_PID]
argsfile [PID_FILE_SLAPD_ARGS]
loglevel 256
...
```

**Step 3** Log in to FusionInsight Manager and choose **Cluster > Services > LdapServer**. Click **More** and select **Restart Service**. In the displayed dialog box, verify the current user identity, and restart the service.

----End

### 7.10.3.12 Updating Encryption Keys of an MRS Cluster

When you install a cluster, the system generates an encryption key automatically to protect the security information stored in the cluster, such as database user passwords and key file access passwords. If you need to change the encryption key for any reason, such as if the original key is compromised or you want to rotate the key periodically, you can update it manually.

#### Impact on the System

- After a cluster key is updated, a new key is generated randomly in the cluster. This key is used to encrypt and decrypt the newly stored data. The old key is not deleted, and it is used to decrypt data encrypted using the old key. After security information is modified, for example, a database user password is changed, the new password is encrypted using the new key.
- When the key is updated, the cluster is stopped and cannot be accessed.

#### Prerequisites

- You have obtained the IP addresses of the active and standby management nodes.
- You have stopped the upper-layer service applications that depend on the cluster.

### Updating an MRS Cluster Key (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **Cluster > Overview > More > Stop** (for MRS 3.3.0 or later, choose **More > Stop** in the upper right corner on the **Homepage**). In the displayed dialog box, enter the password for the current login user.

Click **OK**. Wait until the system displays a message indicating that the cluster is successfully stopped.

**Step 3** Log in to the active management node as user **omm**.

**Step 4** Run the following command to disable logout upon timeout:

```
TMOUT=0
```

 **NOTE**

After the operations in this section are performed, run the **TMOUT=Timeout interval** command to restore the timeout interval in a timely manner. For example, **TMOUT=600** indicates that a user is logged out if the user does not perform any operation within 600 seconds.

**Step 5** Run the following command to switch the directory:

```
cd ${BIGDATA_HOME}/om-server/om/tools
```

**Step 6** Run the following command to update the cluster key:

```
sh updateRootKey.sh
```

Enter **y** as prompted.

The root key update is a critical operation.  
Do you want to continue?(y/n):

The key is updated successfully if the following information is displayed:

```
Step 4-1: The key save path is obtained successfully.
```

```
...
```

```
Step 4-4: The root key is sent successfully.
```

**Step 7** On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and click **Start**. (For MRS 3.3.0 or later, choose **More** > **Start** in the upper right corner of the homepage.)

In the displayed dialog box, click **OK**. Wait until a message is displayed indicating that the startup is successful.

----End

## Updating an MRS Cluster Key (MRS 2.x or Earlier)

**Step 1** Log in to MRS Manager and choose **Services** > **More** > **Stop Cluster**.

In the displayed dialog box, select **I have read the information and understand the impact**. Click **OK**. Wait until the system displays a message indicating that the operation is successful. Click **Finish**. The cluster is stopped successfully.

**Step 2** Log in to the active management node.

**Step 3** Run the following commands to switch the user:

```
sudo su - omm
```

**Step 4** Run the following command to disable logout upon timeout:

```
TMOUT=0
```

**Step 5** Run the following command to switch the directory:

```
cd ${BIGDATA_HOME}/om-0.0.1/tools
```

**Step 6** Run the following command to update the cluster key:

```
sh updateRootKey.sh
```

Enter **y** as prompted.

```
The root key update is a critical operation.
Do you want to continue?(y/n):
```

The key is updated successfully if the following information is displayed:

```
...
Step 4-1: The key save path is obtained successfully.
...
Step 4-4: The root key is sent successfully.
```

**Step 7** On MRS Manager, choose **Services > More > Start Cluster**.

In the displayed dialog box, click **OK**. After **Operation successful** is displayed, click **Finish**. The cluster is started.

----End

### 7.10.3.13 Updating the SSH Key of User omm on MRS Cluster Nodes

#### Scenario

During cluster installation, the system automatically generates the SSH public key and private key for user **omm** to establish the trust relationship between nodes. After the cluster is installed, if the original keys are accidentally disclosed or new keys are used, the system administrator can perform the following operations to manually change the keys.

#### NOTE

This topic is available for MRS 3.x or later.

#### Prerequisites

- The cluster has been stopped.
- No other management operations are being performed.

#### Updating SSH Keys for User omm

**Step 1** Log in as user **omm** to the node whose SSH keys need to be replaced.

If the node is a Manager management node, run the following command on the active management node.

**Step 2** Run the following command to disable logout upon timeout:

```
TMOUT=0
```

 NOTE

After the operations in this section are complete, run the **TMOUT=Timeout interval** command to restore the timeout interval in a timely manner. For example, **TMOUT=600** indicates that a user is logged out if the user does not perform any operation within 600 seconds.

**Step 3** Run the following command to generate a key for the node:

- If the node is a Manager management node, run the following command:  
**sh \${CONTROLLER\_HOME}/sbin/update-ssh-key.sh**
- If the node is a non-Manager management node, run the following command:

**sh \${NODE\_AGENT\_HOME}/bin/update-ssh-key.sh**

If "Succeed to update ssh private key." is displayed when the preceding command is executed, the SSH key is generated successfully.

**Step 4** Run the following command to copy the public key of the node to the active management node:

```
scp ${HOME}/.ssh/id_rsa.pub oms_ip:${HOME}/.ssh/id_rsa.pub_bak
```

*oms\_ip* indicates the IP address of the active management node.

Enter the password of user **omm** to copy the files.

**Step 5** Log in to the active management node as user **omm**.

**Step 6** Run the following command to disable logout on system timeout:

```
TMOUT=0
```

**Step 7** Run the following command to go to the related directory:

```
cd ${HOME}/.ssh
```

**Step 8** Run the following command to add new public keys:

```
cat id_rsa.pub_bak >> authorized_keys
```

**Step 9** Run the following command to move the temporary public key file, for example, /**tmp**.

```
mv -f id_rsa.pub_bak /tmp
```

**Step 10** Copy the **authorized\_keys** file of the active management node to the other nodes in the cluster:


```
scp authorized_keys node_ip:${HOME}/.ssh/authorized_keys
```

*node\_ip* indicates the IP address of another node in the cluster. Multiple IP addresses are not supported.

**Step 11** Run the following command to confirm private key replacement without entering the password:

```
ssh node_ip
```

*node\_ip* indicates the IP address of another node in the cluster. Multiple IP addresses are not supported.

- Step 12** Log in to FusionInsight Manager. On the homepage, click  > **Start** next to the name of the desired cluster to start the cluster. (For MRS 3.3.0 or later, choose **More** > **Start** in the upper right corner of the homepage.)

----End

### 7.10.3.14 Enabling and Disabling Permission Verification on MRS Cluster Components

#### Scenario

HDFS and ZooKeeper verify the permission of users who attempt to access the services in both security and normal clusters by default. Users without related permission cannot access resources in HDFS and ZooKeeper. When the cluster is deployed in normal mode, HBase and YARN do not verify the permission of users who attempt to access the services by default. All users can access resources in HBase and YARN.

Based on actual service requirements, administrators can enable permission verification on HBase and YARN or disable permission verification on HDFS and ZooKeeper in normal clusters.

#### NOTE

This topic is available for MRS 3.x or later.

#### Impact on the System


After the enabling and disabling operations, the service configuration will expire. You need to restart the corresponding service for the configuration to take effect.

#### Enabling Permission Verification on HBase

- Step 1** Log in to FusionInsight Manager.
- Step 2** Click **Cluster**, click the name of the desired cluster, choose **Services** > **HBase**, and click **Configurations**.
- Step 3** Click **All Configurations**.
- Step 4** Search for parameters **hbase.coprocessor.region.classes**, **hbase.coprocessor.master.classes**, and **hbase.coprocessor.regionserver.classes**.

Add the coprocessor parameter **org.apache.hadoop.hbase.security.access.AccessController** to the end of the values of the preceding parameters, and use a comma (,) to separate the values from those of the original coprocessors. The **hbase.coprocessor.region.classes** parameter is used as an example.

**Figure 7-37** hbase.coprocessor.region.classes



**Step 5** Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

## Disabling Permission Verification on HBase

### NOTE

After HBase permission verification is disabled, the existing permission data will be retained. If you want to delete permission information, disable permission verification, enter the HBase shell, and delete table **hbase:acl**.

**Step 1** Log in to FusionInsight Manager.

**Step 2** Click **Cluster**, click the name of the desired cluster, choose **Services > HBase**, and click **Configurations**.

**Step 3** Click **All Configurations**.

**Step 4** Search for parameters **hbase.coprocessor.region.classes**, **hbase.coprocessor.master.classes**, and **hbase.coprocessor.regionserver.classes**.

Delete the coprocessor parameter **org.apache.hadoop.hbase.security.access.AccessController** from the preceding parameter values.

**Step 5** Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

## Disabling Permission Verification on HDFS

**Step 1** Log in to FusionInsight Manager.

**Step 2** Click **Cluster**, click the name of the desired cluster, choose **Services > HDFS**, and click **Configurations**.

**Step 3** Click **All Configurations**.

**Step 4** Search for parameters **dfs.namenode.acls.enabled** and **dfs.permissions.enabled**.

- **dfs.namenode.acls.enabled** indicates whether to enable HDFS ACL. The default value is **true**, indicating that the ACL is enabled. Change the value to **false**.
- **dfs.permissions.enabled** indicates whether to enable permission check for HDFS. The default value is **true**, indicating that permission check is enabled. Change the value to **false**. After the modification, the owner, owner group, and permission of the directories and files in HDFS remain unchanged.

**Step 5** Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

## Enabling Permission Verification on YARN

**Step 1** Log in to FusionInsight Manager.

**Step 2** Click **Cluster**, click the name of the desired cluster, choose **Services > Yarn**, and click **Configurations**.


**Step 3** Click **All Configurations**.

**Step 4** Search for parameter **yarn.acl.enable**.

**yarn.acl.enable** indicates whether to enable the permission check for YARN.

- In normal clusters, the value is set to **false** by default to disable permission check. To enable permission check, change the value to **true**.
- In security clusters, the value is set to **true** by default to enable authentication.

**Figure 7-38** Setting the yarn.acl.enable parameter



**Step 5** Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

## Disabling Permission Verification on ZooKeeper

**Step 1** Log in to FusionInsight Manager.

**Step 2** Click **Cluster**, click the name of the desired cluster, choose **Services > ZooKeeper**, and click **Configurations**.

**Step 3** Click **All Configurations**.

**Step 4** Search for parameter **skipACL**.

**skipACL** indicates whether to skip the ZooKeeper permission check. The default value is **no**, indicating that permission check is enabled. Change the value to **yes**.

**Figure 7-39** Setting the skipACL parameter



**Step 5** Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

## 7.10.3.15 Allowing External Users to Access MRS Clusters in Normal Mode

### Scenario

When the cluster is installed in normal mode, the component clients do not support security authentication and cannot use the **kinit** command. Therefore, nodes outside the cluster cannot use users in the cluster by default. This may result in a user authentication failure when one of these nodes access a component server.



The node administrator can configure a user who has the same name as that of a user for a node outside the cluster, allow the user to log in to the node using the SSH protocol, and connect to the servers of components in the cluster by using the user who logs in to the OS.

 **NOTE**

This topic is available for MRS 3.x or later.

## Prerequisites

- Nodes outside the cluster can connect to the service plane of the cluster.
- The KrbServer service of the cluster is running properly.
- You have obtained the password of user **root** of the node outside the cluster.
- A human-machine user has been planned and added to the cluster, and you have obtained the authentication credential file. For details, see [Creating a User \(MRS 3.x and Later\)](#) and [Downloading MRS Cluster User Credentials](#).

## Allowing External Users to Access Clusters in Normal Mode

**Step 1** Log in to the node where a user is to be added as user **root**.

**Step 2** Run the following command:

```
rpm -qa | grep pam
```

```
rpm -qa| grep krb5-client
```

The following RPM packages are displayed:

```
pam_krb5-32bit-2.3.1-47.12.1
pam-modules-32bit-11-1.22.1
yast2-pam-2.17.3-0.5.211
pam-32bit-1.1.5-0.10.17
pam_mount-32bit-0.47-13.16.1
pam-config-0.79-2.5.58
pam_krb5-2.3.1-47.12.1
pam-doc-1.1.5-0.10.17
pam-modules-11-1.22.1
pam_mount-0.47-13.16.1
pam_ldap-184-147.20
pam-1.1.5-0.10.17
krb5-client-1.6.3
```

**Step 3** Check whether the RPM packages in the list are installed in the OS.

- If yes, go to [Step 5](#).
- If no, go to [Step 4](#).

**Step 4** Obtain the lacked RPM packages from the OS image, upload the files to the current directory, and run the following command to install the RPM package:

```
rpm -ivh *.rpm
```

 **NOTE**

The RPM packages to be installed may bring security risks. The risks that may be brought by the installation of these RPM packages must be taken into consideration during OS hardening.

After the RPM packages are installed, go to [Step 5](#).

**Step 5** Run the following command to configure Kerberos authentication on PAM:

```
pam-config --add --krb5
```

 **NOTE**

If you need to cancel Kerberos authentication and system user login on a non-cluster node, run the **pam-config --delete --krb5** command as user **root**.

**Step 6** Decompress the authentication credential file to obtain **krb5.conf**, use WinSCP to upload this configuration file to the **/etc** directory on the node outside the cluster, and run the following command to configure related permission to enable other users to access the file, such as permission **604**:

```
chmod 604 /etc/krb5.conf
```

**Step 7** Run the following command in the connection session as user **root** to add the corresponding OS user to the human-machine user, and specify **root** as the primary group.

The OS user password is the same as the initial password when the human-machine user is created on Manager.

```
useradd User name -m -d /home/admin_test -g root -s /bin/bash
```

For example, if the name of the human-machine user is **admin\_test**, run the following command:

```
useradd admin_test -m -d /home/admin_test -g root -s /bin/bash
```

 **NOTE**

When you use the newly added OS user to log in to the node by using the SSH protocol for the first time, the system prompts that the password has expired after you enter the user password, and the system prompts that the password needs to be changed after you enter the user password again. You need to enter a new password that meets the password complexity requirements of both the node OS and the cluster.

----End

### 7.10.3.16 Configuring Secure Communication Authorization for an MRS Cluster


MRS clusters provision, manage, and use big data components through the management console. Big data components are deployed in a user's VPC. If the MRS management console needs to directly access big data components deployed in the user's VPC, you need to enable the corresponding security group rules after you have obtained user authorization. This authorization process is called secure communications.

If the secure communications function is not enabled, MRS clusters cannot be created. If you disable the communication after a cluster is created, the cluster status will be **Network channel is not authorized** and the following functions will be affected:

- Installation of cluster components, cluster scaling out or in, upgrading Master node specifications will not be available.
- The cluster running status, alarms, and events cannot be monitored.

- The node management, component management, alarm management, file management, job management, patch management, and tenant management functions on the cluster details page are unavailable.
- The Manager page and the website of each component cannot be accessed.

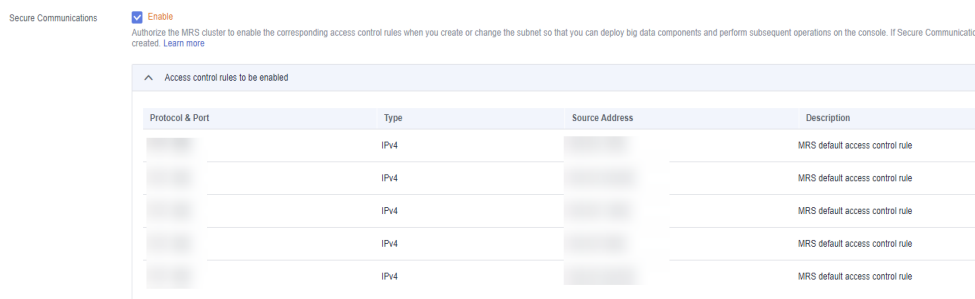
After the secure communications function is enabled again, the cluster status is restored to **Running**, and the preceding functions become available. For details, see [Enabling Secure Communications for Clusters with This Function Disabled](#).

If the security group rules authorized in the cluster are insufficient for you to provision, manage, and use big data components,  is displayed on the right of **Secure Communications**. In this case, click **Update** to update the security group rules. For details, see [Update](#).

## Enabling Secure Communications During Cluster Creation

- Step 1** Log in to the MRS console.
- Step 2** Click **Buy Cluster**. The page for buying a cluster is displayed.
- Step 3** On the displayed page, select **Quick Config**.
- Step 4** Configure cluster information by referring to [Quickly Buying an MRS Cluster](#) or [Manually Buying an MRS Cluster](#).
- Step 5** Select the check box for **Secure Communications**.

**Figure 7-40** Secure communications



- Step 6** Click **Buy Now**.

If Kerberos authentication is enabled for a cluster, check whether Kerberos authentication is required. If yes, click **Continue**. If no, click **Back** to disable Kerberos authentication and then create a cluster.

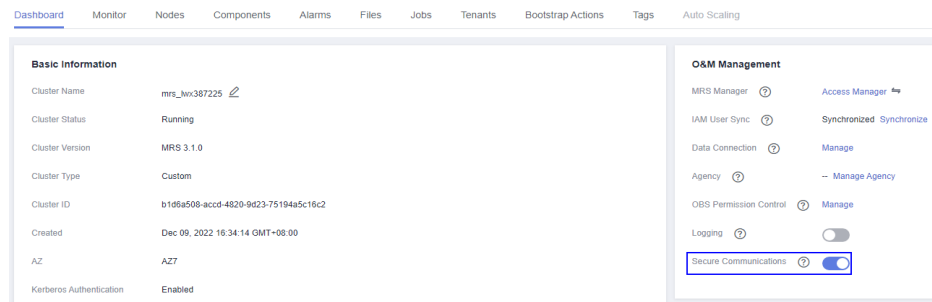
----End

## Disabling Secure Communications After a Cluster Is Created

- Step 1** Log in to the MRS console.
- Step 2** In the active cluster list, click the name of the cluster for which you want to disable secure communications.

The cluster details page is displayed.

**Figure 7-41** Secure Communications



**Step 3** Click the switch on the right of **Secure Communications** to disable authorization. In the dialog box that is displayed, click **OK**.

After the authorization is disabled, the cluster status changes to **Network channel unauthorized**, and some functions of the cluster are unavailable. Exercise caution when performing this operation.

**Figure 7-42** Disabling secure communications

### Disable Secure Communications

If Secure Communications is disabled, the security group rules of the cluster will be deleted. As a result, operations such as this required for O&M cannot be performed on the cluster and some functions of the cluster will be unavailable. Disabling Secure Communications is a high-risk operation. Exercise caution when performing this operation. The following security group rules will be deleted. [Learn more](#)

| Protocol & Port | Type | Source Address | Description                     |
|-----------------|------|----------------|---------------------------------|
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |

- Step 4** If you have enabled critical operation protection (see [Critical Operation Protection](#) on IAM), enter the verification code obtained using the selected verification method to avoid risks and losses caused by misoperations.

**Figure 7-43** Identity verification

Identity Verification

**i** You have enabled operation protection. If you do not require operation protection for critical operations, go to Security Settings > Critical Operations > Operation Protection to disable it. [Disable Identity Verification](#)

Verification Method  SMS  Email  Virtual MFA device [?](#)

Email c\*\*5@huawei.com [Change](#)

Verification Code

----End

## Enabling Secure Communications for Clusters with This Function Disabled

- Step 1** Log in to the MRS console.
- Step 2** In the active cluster list, click the name of the cluster for which you want to enable secure communications.
- The cluster details page is displayed.
- Step 3** Click the switch on the right of **Secure Communications** to enable the function.
- After the function is enabled, the cluster status changes to **Running**.

**Figure 7-44** Enabling secure communications


### Enable Secure Communications

After Secure Communications is enabled, security group rules will be enabled for the cluster. You can perform O&M operations on the cluster and the cluster status is restored to Running. The following lists the security group rules to be enabled. [Learn more](#)

| Protocol & Port | Type | Source Address | Description                     |
|-----------------|------|----------------|---------------------------------|
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |
| TCP : 9022      | IPv4 |                | MRS default security group rule |

----End

## Update

If the security group rules authorized in the cluster are insufficient for you to provision, manage, and use big data components,  is displayed on the right of **Secure Communications**. In this case, click **Update** to update the security group rules. For details, see [Update](#).

**Step 1** Log in to the MRS console.

**Step 2** In the active cluster list, click the name of the cluster for which you want to update secure communications.

The cluster details page is displayed.

**Step 3** Click **Update** on the right of **Secure Communications**.

Figure 7-45 Update



**Step 4** Click **OK**.

----End

## 7.10.4 Changing the Passwords for System Users of an MRS Cluster

### 7.10.4.1 Changing or Resetting the Password for User admin of an MRS Cluster

User **admin** is the system administrator account of Manager. You are advised to periodically change the password by referring to [Changing the Password of User admin](#) to improve system security. If you forget the password, you can reset it by referring to [Resetting the Password of User admin](#).

If the password is changed, the downloaded user credential will be unavailable. Download the authentication credential again, and replace the old one.

### Changing the Password of User admin

#### Changing the password of user admin on Manager:

You can change the password for the user **admin** on Manager for clusters with Kerberos authentication enabled or clusters with Kerberos authentication disabled but with the EIP function enabled.

**Step 1** Log in to Manager as user **admin**.

- For MRS 2.x or earlier, click the username in the upper right corner of the page and choose **Change Password**.
- For MRS 3.x or later, hover over **Hello, admin** in the upper right corner of the page and choose **Change Password**.

**Step 2** On the **Change Password** page, set **Old Password**, **New Password**, and **Confirm Password**.

 NOTE

The default password complexity requirements are as follows:

- For MRS 2.x or earlier:
  - The password must contain at least eight characters.
  - The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#%&\*()-\_+=\|[];:;";<.>/?').
  - The password cannot be the username or the reverse username.
- MRS 3.x or later:
  - The password must contain at least eight characters.
  - The password must contain at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (~!?.;:\_'(){}[]/<>@#% ^&\*+|\=).
  - The password cannot be the same as the username or the username spelled backwards.
  - The password cannot be a common easily-cracked password.
  - The password cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in [Configuring Password Policies for MRS Cluster Users](#).

**Step 3** Click **OK**. Log in to Manager with the new password.

----End

**Changing the password of user admin on the cluster node** (versions later than MRS 2.x):

**Step 1** Update the client of the active management node. For details, see [Updating the MRS Cluster Client After the Server Configuration Expires](#).

**Step 2** Log in to the active management node.

**Step 3** (Optional) To change the password as user **omm**, run the following command to switch the user:

```
sudo su - omm
```

**Step 4** Run the following command to switch to the client directory, for example, **/opt/client**.

```
cd /opt/client
```

**Step 5** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 6** Run the following command to change the password of user **admin**: This operation takes effect in the whole cluster.

```
kpasswd admin
```

Enter the old password and then enter a new password twice.

For the cluster, the default password complexity requirements are as follows:

- The password must contain at least eight characters.



- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#\$%^&\*()-\_+=\|[]{};:'''<.>/?).
- The password cannot be the username or the reverse username.

----End

## Resetting the Password of User admin

**Step 1** Log in to the **Master1** node.

**Step 2** (Optional) To change the password as user **omm**, run the following command to switch the user:

```
sudo su - omm
```

**Step 3** Go to the client directory.

```
cd Client installation directory
```

**Step 4** Run the following command to set environment variables:

```
source bigdata_env
```

**Step 5** Run the following command to log in to the console as user **kadmin/admin**:

```
kadmin -p kadmin/admin
```

### NOTE

The default password of user **kadmin/admin** is **KAdmin@123** for MRS 2.x or earlier and **Admin@123** for MRS 3.x or later. The system displays a message indicating that the password has expired upon the first login. Change the password as prompted. Keep the password secure as it cannot be retrieved once lost.

**Step 6** Run the following command to reset the password of user **admin**:

```
cpw admin
```

For the cluster, the default password complexity requirements are as follows:

- For MRS 2.x or earlier:
  - The password must contain at least eight characters.
  - The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#\$%^&\*()-\_+=\|[]{};:'''<.>/?).
  - The password cannot be the username or the reverse username.
- MRS 3.x or later:
  - The password must contain at least eight characters.
  - The password must contain at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (~`!?,;,-\_'(){}[]/<>@#\$\$%^&\*+|\=).
  - The password cannot be the same as the username or the username spelled backwards.
  - The password cannot be a common easily-cracked password.

- The password cannot be the same as the password used in the last  $N$  times.  $N$  indicates the value of **Repetition Rule** in [Configuring Password Policies for MRS Cluster Users](#).

----End

### 7.10.4.2 Changing the Passwords for OS Users of an MRS Cluster Node

This section describes how to periodically change the login passwords for the OS users **omm**, **ommdba**, and **root** on MRS cluster nodes to improve system O&M security.

You do not need to set a unified password for the OS users on each node.

#### Prerequisites

- You have obtained the IP address of the node where the passwords of users **omm** and **ommdba** are to be changed.
- You have obtained the password of user **root** before changing the passwords of users **omm** and **ommdba**.

### Changing the Password of the OS User of an MRS Cluster Node

**Step 1** Log in to the node that requires password change as user **root**.

**Step 2** Run the following command to switch the user:

```
sudo su - root
```

**Step 3** Run the following command to change the passwords of users **omm**, **ommdba**, or **root**:

```
passwd omm
```

```
passwd ommdba
```

```
passwd root
```

For example, after you run the command to change the password of user **omm**, the following information is displayed:

```
Changing password for user omm.
New password:
Retype new password:
```

Enter the new password and confirm the password. The password change policies for an OS vary according to the OS that is used.

#### NOTE

The default password complexity requirements of the MRS cluster are as follows:

- The password must contain at least eight characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#\$\$%^&\*()-\_+=\| [{}];:~",<.>/?').
- The new password cannot be the same as last five historical passwords.

----End

### 7.10.4.3 Changing the Password for the Kerberos Administrator of an MRS Cluster

This section describes how to periodically change the password for the Kerberos or OMS Kerberos administrator **kadmin** of an MRS cluster (MRS 3.x or later) to improve system O&M security.

If the password is changed, the downloaded user credential will be unavailable. Redownload the authentication credential and replace the old one.

#### Prerequisites

- For MRS 2.x or earlier, the client has been installed on the **Master1** node.
- For MRS 3.x or later, the client has been installed on any node in the cluster and the IP address of the node has been obtained.

#### Changing the Password of the Kerberos Administrator

If the current MRS version is 3.x or later, changing the password of this user will change the password of the OMS Kerberos administrator.

**Step 1** Log in to a cluster node.

- For MRS 3.x or later, log in to the node where the client is installed as user **root** using the node IP address.
- For MRS 2.x or earlier: Log in to the **Master1** node.

**Step 2** (Optional) To change the password as user **omm**, run the following command to switch to user **omm**:

```
sudo su - omm
```

**Step 3** Run the following command to go to the client directory, for example, **/opt/hadoopclient**:

```
cd /opt/hadoopclient
```

**Step 4** Run the following command to set environment variables:

```
source bigdata_env
```

**Step 5** Run the following command to change the password of **kadmin/admin**. This operation takes effect for all servers. Keep the password secure because it cannot be retrieved once lost.

```
kpasswd kadmin/admin
```

Enter the password (default: **Admin@123**) and set a new password. The new password must meet the following complexity requirements:

- For MRS 2.x or earlier:
  - The password must contain at least eight characters.
  - The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters (`'~!@#$$%^&*()-_+=\|[{]};:","<.>/?`).
  - The password cannot be the username or the reverse username.

- MRS 3.x or later:
  - The password must contain at least eight characters.
  - The password must contain at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (`~`!?,;,-_'(){}[]/<>@#$$%^&*+|\=`).
  - The password cannot be the same as the username or the username spelled backwards.
  - The password cannot be a common easily-cracked password.
  - The password cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in [Configuring Password Policies for MRS Cluster Users](#).

----End

## Changing the Password of the OMS Kerberos Administrator

This operation is supported only in MRS 3.x or later.

Changing this user's password will also update the Kerberos administrator password.

**Step 1** Log in to any management node in the cluster as user **omm**.

**Step 2** Run the following command to go to the directory:

```
cd ${BIGDATA_HOME}/om-server/om/meta-0.0.1-SNAPSHOT/kerberos/scripts
```

**Step 3** Run the following command to set environment variables:

```
source component_env
```

**Step 4** Run the following command to change the password of **kadmin/admin**. This operation takes effect for all servers. Keep the password secure because it cannot be retrieved once lost.

```
kpasswd kadmin/admin
```

Enter the user password and set a new password. The new password must meet the following complexity requirements:

- The password contains at least 8 characters.
- The password contains at least four types of the following: uppercase letters, lowercase letters, numbers, and special characters (`~`!?,;,-_'(){}[]/<>@#$$%^&*+|\=`).
- The password cannot be the same as the username or the username spelled backwards.
- The password cannot be a common easily-cracked password.
- The password cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in [Configuring Password Policies for MRS Cluster Users](#).

----End

### 7.10.4.4 Changing the Passwords for Manager Users of an MRS Cluster

Passwords of human-machine system users must be regularly changed to ensure MRS cluster security. This section describes how to change passwords on Manager.

If you have the permission to use Manager, you can change the passwords on Manager.

If you do not have the permission to use Manager, you can change the passwords on the cluster client.

#### Impact on the System

If you have downloaded a user authentication file, download it again and obtain the keytab file after changing the password of the MRS cluster user.

#### Prerequisites

- You have obtained the current password policy. For details, see [Configuring Password Policies for MRS Cluster Users](#).
- If you need to use a client to reset passwords, ensure that the cluster client has been installed on any node in the cluster.

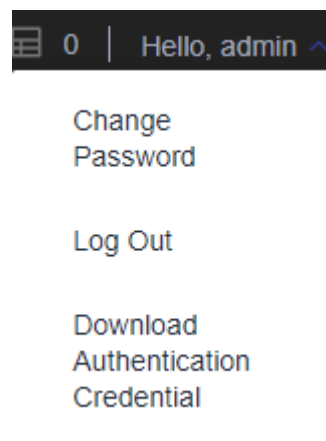
#### Changing Passwords on Manager

**Step 1** Log in to Manager as a user who has the user management permission, for example, **admin**.

**Step 2** Hover over the username in the upper right corner of the page.

Choose **Change Password**.

**Figure 7-46** Choosing Change Password



**Step 3** On the displayed page, set **Current Password**, **New Password**, and **Confirm Password**, and click **OK**.

The default password complexity requirements are as follows:

- For MRS 2.x or earlier:
  - The password must contain at least eight characters.

- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#\$\$%^&\*()-\_+=\|[{]}:;'"<.>/?').
- The password cannot be the username or the reverse username.
- MRS 3.x or later:
  - The password must contain at least eight characters.
  - The password must contain at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (~`!?,;,-\_'(){}[]/<>@#\$\$%^&\*+|\=).
  - The password cannot be the same as the username or the username spelled backwards.
  - The password cannot be a common easily-cracked password.
  - The password cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in [Configuring Password Policies for MRS Cluster Users](#).

----End

## Changing Passwords on the Client

This function is only available in MRS 3.x or later.

**Step 1** Log in to the node where the client is installed as the client installation user.

**Step 2** Run the following command to switch to the client directory, for example, **/opt/client**:

```
cd /opt/client
```

**Step 3** Run the following command to set environment variables:

```
source bigdata_env
```

**Step 4** Change the user password.

```
kpasswd System username
```

For example, change the password of the system user **test1**.

```
kpasswd test1
```

Enter the old password as prompted and set a new password. The default password complexity requirements are as follows:

- The password must contain at least eight characters.
- The password must contain at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (~`!?,;,-\_'(){}[]/<>@#\$\$%^&\*+|\=).
- The password cannot be the same as the username or the username spelled backwards.
- The password cannot be a common easily-cracked password.
- The password cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in [Configuring Password Policies for MRS Cluster Users](#).

 NOTE

If an error occurs during the running of the **kpasswd** command, try the following operations:

- Close the SSH session and reopen it.
- Run **kdestroy** and then **kpasswd**.

----End

### 7.10.4.5 Changing the Password for a Regular LDAP User of an MRS Cluster

This section describes how to periodically change the passwords for LDAP administrators and users of an MRS cluster to improve system O&M security.

- Usernames in MRS 3.1.0:
  - LDAP administrators: `cn=root, dc=hadoop, dc=com`
  - LDAP users: `cn=pg_search_dn, ou=Users, dc=hadoop, dc=com`
- Usernames in MRS 2.x or earlier:
  - LDAP administrators: `rootdn:cn=root, dc=hadoop, dc=com`
  - LDAP users: `pg_search_dn:cn=pg_search_dn, ou=Users, dc=hadoop, dc=com`

 NOTE

- This section only applies to MRS 2.x or earlier and MRS 3.1.0. For versions later than MRS 3.1.0, refer to [Modifying the OMS Service Configuration](#).
- For MRS 3.1.0 clusters:
  - Changing these user passwords will also update OMS LDAP administrator or user passwords.
  - If the cluster is upgraded from an earlier version, LDAP administrator passwords will inherit the password policy of the old cluster. To ensure system security, you are advised to change the passwords after the cluster upgrade.

### Impact on the System

- For MRS 2.x or earlier, all services need to be restarted after passwords are changed, during which the services are unavailable.
- For MRS 3.1.0 clusters:
  - Changing the user password of the `LdapServer` service is a high-risk operation and requires restarting the `KrbServer` and `LdapServer` services. If `KrbServer` is restarted, users may fail to be queried by running the **id** command on nodes in the cluster temporarily. Therefore, exercise caution when restarting `KrbServer`.
  - After the passwords for LDAP users **cn=pg\_search\_dn, ou=Users, dc=hadoop**, and **dc=com** are changed, the users may be locked in the LDAP component. You are advised to unlock them after password change. For details, see [Unlocking the LDAP Management Account of the MRS Cluster](#).

## Prerequisites (MRS 3.1.0)

Before changing the passwords for LDAP users **cn=pg\_search\_dn**, **ou=Users**, **dc=hadoop**, and **dc=com**, ensure that the users are not locked by running the following command on the active management node of the cluster:

```
ldapsearch -H ldaps://Floating IP address of OMS:OLDAP port-LLL -x -D
cn=pg_search_dn,ou=Users,dc=hadoop,dc=com -W -b
cn=pg_search_dn,ou=Users,dc=hadoop,dc=com -e ppolicy
```

Enter the password for the LDAP user **pg\_search\_dn**. If the following information is displayed, the user is locked. In this case, unlock the user by referring to [Unlocking the LDAP Management Account of the MRS Cluster](#).

```
ldap_bind: Invalid credentials (49); Account locked
```

### NOTE

- To obtain the OLDAP port, log in to FusionInsight Manager, choose **System** > **OMS** > **oldap** > **Modify Configuration**, and view the value of **LDAP Listening Port**.
- The password for the LDAP user **pg\_search\_dn** is randomly generated by the system, which can be from the `/etc/sss/sss.conf` or `/etc/ldap.conf` file on the active node.

## Changing the Password for a Regular LDAP User of an MRS Cluster

For MRS 3.1.0:

**Step 1** Log in to FusionInsight Manager and choose **Cluster** > **Services** > **LdapServer**.

**Step 2** In the upper right corner of the **Overview** page, choose **More** > **Change Database Password**. In the displayed dialog box, enter the password of the current login user and click **OK**.

**Step 3** In the **Change Password** dialog box, select the user whose password needs to be changed in the **User Information** drop-down list.

**Step 4** Enter the old password in the **Old Password** text box, and enter the new password in the **New Password** and **Confirm Password** text boxes.

The default password complexity requirements are as follows:

- The password must contain 16 to 32 characters.
- The password contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (`~!@#%&*( )-_=+|[]{};,<.>/?`).
- The password cannot be the same as the username or the username spelled backwards.
- The password cannot be the same as the current password.

**Step 5** Select **I have read the information and understood the impact** and click **OK** to confirm the modification and restart the service.

----End

For MRS 2.x or earlier:

**Step 1** On MRS Manager, choose **Services** > **LdapServer** > **More**.



- Step 2** Click **Change Password**. In the displayed dialog box, enter the old password and click **OK**.
- Step 3** In the **Change Password** dialog box, select the user whose password needs to be modified in the **User Information** drop-down box.
- Step 4** Enter the old password in the **Old Password** text box, and enter the new password in the **New Password** and **Confirm Password** text boxes.

The default password complexity requirements are as follows:

- The password contains 16 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters (`~!@#%&*(-_+=\| [{}];:;<.>/?)`).
- The password cannot be the username or the reverse username.
- The new password cannot be the same as the current password.

 **NOTE**

The default password of the LDAP administrator `rootdn:cn=root,dc=hadoop,dc=com` is `LdapChangeMe@123`, and that of the LDAP user `pg_search_dn:cn=pg_search_dn,ou=Users,dc=hadoop,dc=com` is `pg_search_dn@123`. Periodically change the passwords and keep them secure.

- Step 5** Select **I have read the information and understand the impact**, and click **OK** to confirm the modification and restart the service.

----End

### 7.10.4.6 Changing the LDAP Administrator Password for an MRS Cluster

It is recommended that the administrator periodically changes the passwords of LDAP administrator accounts `cn=krbkdc,ou=Users,dc=hadoop,dc=com` and `cn=krbadmin,ou=Users,dc=hadoop,dc=com` to improve the system O&M security.

 **NOTE**

This section applies only to MRS 3.1.0. For versions later than MRS 3.1.0, see [Modifying the OMS Service Configuration](#).

### Impact on the System

- You need to restart the KrbServer service after changing the password.
- After the password is changed, check whether the LDAP administrator accounts `cn=krbkdc,ou=Users,dc=hadoop,dc=com` and `cn=krbadmin,ou=Users,dc=hadoop,dc=com` are locked, run the following command on the active management node of the cluster to check whether `krbkdc` is locked (the method for user `krbadmin` is similar):

 **NOTE**

OLdap port number obtaining method:

1. Log in to FusionInsight Manager, choose **System > OMS > oldap > Modify Configuration**:
2. The **LDAP Listening Port** parameter value is **oldap port**.

```
ldapsearch -H ldaps://OMS_FLOAT_IP address:OLdap port -LLL -x -D
cn=krbkdc,ou=Users,dc=hadoop,dc=com -W -b
cn=krbkdc,ou=Users,dc=hadoop,dc=com -e ppolicy
```

Enter the password of the LDAP administrator account **krbkdc**. The default password is **LdapChangeMe@123**. If the following message is displayed, the account is locked. For details about how to unlock the account, see [Unlocking the LDAP Management Account of the MRS Cluster](#).

```
ldap_bind: Invalid credentials (49); Account locked
```

## Prerequisites

You have obtained the management node IP address.

## Changing the Password of the LDAP Administrator

**Step 1** Log in to the active management node as user **omm** with the IP address of the active management node.

**Step 2** Run the following command to go to the related directory:

```
cd ${BIGDATA_HOME}/om-server/om/meta-0.0.1-SNAPSHOT/kerberos/scripts
```

**Step 3** Run the following command to change the password of the LDAP administrator account:

```
./okerberos_modpwd.sh
```

Enter the old password and then enter a new password twice.

The password must meet the following complexity requirements:

- Contains 16 to 32 characters.
- Contains at least three types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (`~!@#$$%^&*()-_+=|[]{};,<.>/?`).
- Cannot be the same as the current password.

If the following information is displayed, the password is changed.

```
Modify kerberos server password successfully.
```

**Step 4** Log in to FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and choose **Services > KrbServer**. On the displayed page, choose **More > Restart Service**.

Enter the password and do not select **Restart upper-layer services**. Click **OK** to restart the KrbServer service.

----End

### 7.10.4.7 Changing the Passwords for MRS Cluster Component Running Users

You are advised to regularly change the passwords for MRS cluster component running users to enhance system O&M security.

- For MRS 2.x or earlier:

- If the initial password is randomly generated by the system, reset the password.
- If the password is changed, the downloaded user credential will be unavailable. Redownload the authentication credential and replace the old one.
- In MRS 3.x or later, component running users are classified into the following types based on whether their initial passwords are randomly generated by the system:
  - If the initial password of a component running user is randomly generated by the system, the user is of the machine-machine type.
  - If the initial password of a component running user is not randomly generated by the system, the user is of the human-machine type.

## Impact on the System

For MRS 3.x or later, the initial password for the component running user is randomly generated by the system. After changing the password, the MRS cluster needs to be restarted, which may cause temporary interruption of services during the restart.

## Prerequisites

- For MRS 2.x or earlier, the client has been installed on the **Master1** node.
- For MRS 3.x or later, the client has been installed on any node in the cluster and the IP address of the node has been obtained.

## Changing the Passwords for MRS Cluster Component Running Users (MRS 3.x or Later)

**Step 1** Log in to the node where the client is installed as the client installation user.

**Step 2** Go to the client directory.

```
cd Client installation directory
```

**Step 3** Run the following command to set environment variables:

```
source bigdata_env
```

**Step 4** Run the following command and enter the password for user **kadmin/admin** to log in to the kadmin console:

```
kadmin -p kadmin/admin
```

### NOTE

The default password for the user **kadmin/admin** is **KAdmin@123**, which will expire upon your first login. Change the password as prompted. Keep the password secure as it cannot be retrieved once lost.

**Step 5** Run the following command to change the password for an internal component running user:

```
cpw Internal component username
```

Example: **cpw hdfs**

The username **hdfs** is given as an example. Replace it with the actual one.

The default password complexity requirements are as follows:

- The password contains at least 8 characters.
- The password must contain at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (`~`!?,;-'(){}[]/<>@#$$%^&*+|\=`).
- The password cannot be the same as the username or the username spelled backwards.
- The password cannot be a common easily-cracked password, for example, **Admin@12345**.
- The password cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in [Configuring Password Policies for MRS Cluster Users](#). This policy only applies to human-machine accounts.

 **NOTE**

Run the following command to view user information:

```
getprinc Internal system username
```

Example: `getprinc hdfs`

**Step 6** Determine the type of the user whose password needs to be changed.

- If the user is a machine-machine user, go to [Step 7](#).
- If the user is a human-machine user, the password is successfully changed and no further action is required.

**Step 7** Log in to FusionInsight Manager.

**Step 8** On the home page, click  or **More** and click **Restart**.

**Step 9** In the displayed dialog box, enter the password for the current login user and click **OK**.

**Step 10** In the displayed restart confirmation dialog box, click **OK**.

**Step 11** Wait until the system displays a message indicating that the restart is successful.

----End

## Changing the Passwords for MRS Cluster Component Running Users (MRS 2.x or Earlier)

**Step 1** Log in to the **Master1** node.

**Step 2** (Optional) To change the password as user **omm**, run the following command to switch the user:

```
sudo su - omm
```

**Step 3** Run the following command to switch to the client directory, for example, **/opt/client**:

```
cd /opt/client
```

**Step 4** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 5** Run the following command to log in to the console as user **kadmin/admin**:

```
kadmin -p kadmin/admin
```

 **NOTE**

The default password of user **kadmin/admin** is **KAdmin@123**, which will expire upon your first login. Change the password as prompted. Keep the password secure because it cannot be retrieved once lost.

**Step 6** Run the following command to reset the password of a component running user. This operation takes effect for all servers.

```
cpw Component running user name
```

For example, to reset the password of user admin, run the **cpw admin** command.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#\$%^&\*()-\_+=\|[{]};:","<.>/?').
- The password cannot be the username or the reverse username.

----End

## 7.10.5 Changing the Passwords for Database Users of an MRS Cluster

### 7.10.5.1 Changing the Password for the OMS Database Administrator

You are advised to regularly change the password for the OMS database administrator to enhance system O&M security.

**Step 1** Log in to the active management node as user **root**.

 **NOTE**

To ensure smooth cluster running, you are not allowed to change the password for the user **ommdba** on the standby management node. Change the password on the active management node only.

**Step 2** Run the following command to switch the user:

```
sudo su - omm
```

**Step 3** Run the following command to switch the directory:

```
cd $OMS_RUN_PATH/tools
```

**Step 4** Change the password for the user **ommdba**.

```
mod_db_passwd ommdba
```

**Step 5** Enter the old password of user **ommdba** and enter a new password twice.

The password complexity requirements are as follows:

- The password contains 16 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters ('~!@#\$\$%^&\*()-\_+=\| [{}];:~",<.>/?).
- The password cannot be the username or the reverse username.
- The password cannot be the same as the last 20 historical passwords.

If the following information is displayed, the password is changed successfully.

```
Congratulations, update [ommdba] password successfully.
```

----End

### 7.10.5.2 Changing the Password for an OMS Database Access User

This section describes how to regularly change the password for an OMS database access user to enhance system O&M security.

#### Impact on the System

The OMS service needs to be restarted for the new password to take effect. The service is unavailable during the restart.

### Changing the Password for an OMS Database Access User (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager and choose **System > OMS > gaussDB > Change Password**.

**Step 2** Locate the row that contains the user **omm** and click **Change Password** in the **Operation** column.

**Step 3** In the displayed dialog box, enter the password for the current login user and click **OK**.

**Step 4** Enter the old and new passwords as prompted.

The password must meet the following complexity requirements:

- The password must contain 8 to 32 characters.
- The password contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!@#\$\$%^&\*()-+\_=\| [{}];:~",<.>/?).
- The password cannot be the same as the username or the username spelled backwards.
- The password cannot be the same as the last 20 historical passwords.

**Step 5** Click **OK**. Wait until the system displays a message indicating that the operation is successful.

**Step 6** Locate the row that contains the user **omm** and click **Restart OMS Service** in the **Operation** column.

**Step 7** In the displayed dialog box, enter the password for the current login user and click **OK**.

**Step 8** In the displayed confirmation dialog box, click **OK**.

----End

## Changing the Password for an OMS Database Access User (MRS 2.x or Earlier)

**Step 1** Log in to MRS Manager and click **System**.

**Step 2** In the **Permission** area, click **Change OMS Database Password**.

**Step 3** Locate the row that contains the user **omm** and click **Change Password** in the **Operation** column.

The password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters ('~!@#\$%^&\*()-\_+=\|[{]}:;";<.>/?').
- The password cannot be the username or the reverse username.
- The password cannot be the same as the last 20 historical passwords.

**Step 4** Click **OK**. When **Operation successful** is displayed, click **Finish**.

**Step 5** Locate the row that contains the user **omm** and click **Restart OMS Service** in the **Operation** column.

### NOTE

If the password is changed but the OMS database is not restarted, the status of the user **omm** changes to **Waiting to restart** and the password cannot be changed until the OMS database is restarted.

**Step 6** In the displayed dialog box, select **I have read the information and understand the impact**. Click **OK**, and restart the OMS service.

----End

### 7.10.5.3 Changing the Passwords for Database Users of MRS Cluster Components

This section describes how to regularly change the passwords for database users of cluster components to enhance system O&M security.

#### Impact on the System

The services need to be restarted for the new password to take effect. The services are unavailable during the restart.

## Changing the Passwords for Database Users of MRS Cluster Components (MRS 3.1.0)

This function is only available in MRS 3.1.0. For later versions, refer to [Resetting the MRS Component Database User Password](#).

**Step 1** Log in to FusionInsight Manager and choose **Cluster > Services**.

**Step 2** Click the name of the service whose database user password is to be reset. On the displayed **Dashboard** page, click **Stop Service**.

In the displayed dialog box, enter the password for the current login user and click **OK**.

Once you have confirmed the impact of stopping the service, wait for it to come to a complete stop.

**Step 3** Click the service whose database user password is to be changed, click **More**, and select **Change Database Password**. On the displayed page, enter the password for the current login user and click **OK**.

**Step 4** Enter the old and new passwords as prompted.

The password must meet the following complexity requirements:

- The database user password contains 8 to 32 characters.
- The password contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!@#\$%^&\*()-+\_=|[]{};":',<.>/?).
- The password cannot be the same as the username or the username spelled backwards.
- The password cannot be the same as the last 20 historical passwords.

**Step 5** Select **I have read the information and understand the impact** and click **OK**.

**Step 6** After the password is changed, choose **More > Restart Service**. In the displayed dialog box, enter the password for the current login user, click **OK**, select **Restart the upper-layer services**, and click **OK**.

----End

## Changing the Passwords for Database Users of MRS Cluster Components (MRS 2.x or Earlier)

**Step 1** On MRS Manager, click **Services** and click the name of the database user service to be modified.

**Step 2** Determine the component database user whose password is to be changed.

- To change the password of the DBService database user, go to [Step 3](#).
- To change the password for the Hive, Hue, or Loader database user, click **Stop Service** and go to [Step 3](#).

**Step 3** Choose **More > Change Password**.

**Step 4** Enter the old and new passwords as prompted.

The password complexity requirements are as follows:



- The password of the DBService database user contains 16 to 32 characters. The password of the Loader, Hive, or Hue database user contains 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters ('~!@#\$%^&\*()-\_+=+|[{}];:":'<.>/?).
- The password cannot be the username or the reverse username.
- The password cannot be the same as the last 20 historical passwords.

**Step 5** Click **OK**. The system automatically restarts the corresponding service. When **Operation successful** is displayed, click **Finish**.

----End

### 7.10.5.4 Resetting the MRS Component Database User Password

MRS cluster components use random default passwords for connecting to the DBService database. For operations security, you need to periodically reset the passwords of component database users.

#### NOTE

This section applies only to MRS 3.1.2 or later. For versions earlier than MRS 3.1.2, see [Changing the Passwords for Database Users of MRS Cluster Components](#).

## Impact on the System

To reset passwords, you need to stop and then restart services, during which services are unavailable.

## Resetting the Component Database User Password

**Step 1** Log in to FusionInsight Manager and choose **Cluster > Services**.

**Step 2** Click the name of the service whose database user password is to be reset, for example, **Kafka**, and click **Stop Service** on the **Dashboard** page.

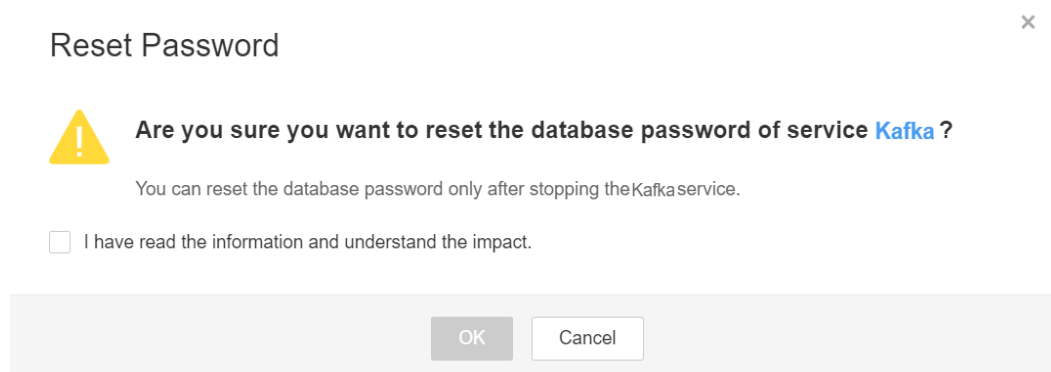
In the displayed dialog box, enter the password of the current login user and click **OK**.

After confirming the impact of stopping the service, wait until the service is stopped.

**Step 3** On the **Dashboard** page, choose **More > Reset Database Password**.

In the displayed dialog box, enter the password of the current login user and click **OK**.

Select "I have read the information and understand the impact", and click **OK**.

**Figure 7-47** Selecting the confirmation check box

**Step 4** After the password is reset, click **Start Service** on the **Dashboard** page.

**Step 5** In the displayed dialog box, click **OK** and wait until the service is started.

----End

### 7.10.5.5 Resetting the Password for User omm in DBService

#### Scenario

The default password of the DBService database user **omm** in the MRS cluster is randomly generated. You need to periodically reset the password to improve system O&M security.

#### NOTE

This topic is available for MRS 3.2.0-LTS.1 and later versions only. Versions earlier than MRS 3.2.0-LTS.1 do not support password resetting.

#### Resetting the Password

**Step 1** Log in to FusionInsight Manager and choose **Cluster > Services > DBService**.

**Step 2** On the **Dashboard** page, choose **More > Reset Database Password**.

In the displayed dialog box, enter the password of the current login user and click **OK**.

Select **I have read the information and understand the impact**, and click **OK**.

**Step 3** After the password is reset, click **More** and select **Service Restart Rolling** on the **Dashboard** page.

In the displayed dialog box, enter the password of the current login user and click **OK**.

**Step 4** Confirm the impact of restarting the service, click **OK**, and wait until the service is started.

----End

### 7.10.5.6 Changing the Password for User compdbuser of the DBService Database

You need to regularly change the password of the OMS database to enhance system security and maintenance.

 **NOTE**

This topic is available for MRS 3.x and later versions only.

**Step 1** Log in to FusionInsight Manager, choose **Cluster > Services > DBService**, click **Instance**, and view the IP address of the active DBService node.

**Step 2** Log in to the active DBService node as user **root**.

 **NOTE**

The password of user **compuserdb** cannot be changed on the standby DBService node. Change the password on the active management node only.

**Step 3** Switch to the **\$DBSERVER\_HOME** directory and configure environment variables.

```
su - omm
```

```
cd $DBSERVER_HOME
```

```
source .dbservice_profile
```

**Step 4** Change the password of user **compdbuser** as user **omm** of the DBService database.

```
gsql -U omm -W omm Password of user omm of the DBService database -d postgres -p 20051 -c "alter user compdbuser identified by 'New password valid until 'Expiration time';"
```

 **NOTE**

- For the initial password of user **omm** in the DBService database, see [MRS Cluster User Accounts](#).
- The new password must meet the following complexity requirements:
  - Contains 16 to 32 characters.
  - Contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!@#\$\$%^&\*()-+\_=|[{}];:","<>/?).
  - Cannot be the same as the username or the username spelled backwards.
  - Cannot be the same as the last 20 historical passwords.
- The expiration time format is xxxx-xx-xx, for example, **2020-10-31**.

If the following information is displayed, the modification is successful:

```
ALTER ROLE
```

```
----End
```

## 7.11 Viewing and Configuring MRS Alarm Events

## 7.11.1 Viewing MRS Cluster Events

The event list displays information about all events in a cluster, such as service restart and service termination.

Events are listed in the event list in chronological order by default, with the most recent events displayed at the top.

### Prerequisites

- The IAM users have been synchronized in advance. You can do this by clicking **Synchronize** next to **IAM User Sync** on the **Dashboard** page of the cluster details.
- You have logged in to MRS Manager. For how to log in, see [Accessing MRS Manager](#).

### Viewing Cluster Events on the Management Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** Choose **Alarms > Events**. On the displayed page, view event information.

**Step 4** Click **Export All**. In the displayed dialog box, select a save type and click **OK**.

----End



### Viewing Cluster Events on Manager


**For MRS 3.x or later:**

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Alarm > Events**. On the displayed **Events** page, you can view information about all events in the cluster, including the event name, ID, severity, generation time, object, and location. By default, the latest 10 events are displayed on each page.

#### NOTE

- Click **Export All** to export all event details.
- You can click  to manually refresh the current page and click  to filter columns to display.
- You can filter events by object or cluster.
- You can click **Advanced Search** to search for events by event ID, name, severity, start time, or end time.

You can click  on the left of an event to view detailed event parameters. [Table 7-85](#) describes the parameters.

**Table 7-85** Event parameters


| Parameter       | Description                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Event ID        | Event ID.                                                                                                                                                                                                                                                                                                                                                                                         |
| Event Name      | Event name.                                                                                                                                                                                                                                                                                                                                                                                       |
| Event Severity  | Event severity. The options are <b>Critical</b> , <b>Major</b> , <b>Minor</b> , and <b>Suggestion</b> .                                                                                                                                                                                                                                                                                           |
| Generated       | Time when an event is generated.                                                                                                                                                                                                                                                                                                                                                                  |
| Object          | Possible cause of an event.                                                                                                                                                                                                                                                                                                                                                                       |
| Serial Number   | Number of events generated by the system.                                                                                                                                                                                                                                                                                                                                                         |
| Location        | Detailed information for locating the event, which includes the following: <ul style="list-style-type: none"><li>• <b>Source</b>: cluster for which the event is generated.</li><li>• <b>ServiceName</b>: service for which the event is generated.</li><li>• <b>RoleName</b>: role for which the event is generated.</li><li>• <b>HostName</b>: host for which the event is generated.</li></ul> |
| Additional Info | Error information.                                                                                                                                                                                                                                                                                                                                                                                |
| Event Cause     | Possible cause of an event.                                                                                                                                                                                                                                                                                                                                                                       |
| Source          | Cluster name.                                                                                                                                                                                                                                                                                                                                                                                     |

----End

**For MRS 2.x or earlier:**

**Step 1** Log in to MRS Manager.

**Step 2** Choose **Alarms > Events**. On the displayed page, you can view information about all events in the cluster, including the event name, ID, severity, and generation time.

You can click  on the left of an event to view detailed event parameters. [Table 7-86](#) describes the parameters.

**Table 7-86** Event parameters

| Parameter      | Description                                                                                             |
|----------------|---------------------------------------------------------------------------------------------------------|
| Event ID       | Specifies the ID of an event.                                                                           |
| Event Severity | Event severity. The options are <b>Critical</b> , <b>Major</b> , <b>Minor</b> , and <b>Suggestion</b> . |
| Event Name     | Name of the generated event.                                                                            |

| Parameter | Description                                  |
|-----------|----------------------------------------------|
| Generated | Time when the event is generated.            |
| Location  | Detailed information for locating the event. |

----End

## Common Events of an MRS Cluster

**Table 7-87** Common events

| Event ID | Event Name                                                      |
|----------|-----------------------------------------------------------------|
| 12019    | Stop Service                                                    |
| 12020    | Delete Service                                                  |
| 12021    | Stop RoleInstance                                               |
| 12022    | Delete RoleInstance                                             |
| 12023    | Delete Node                                                     |
| 12024    | Restart Service                                                 |
| 12025    | Restart RoleInstance                                            |
| 12026    | Manager Switchover                                              |
| 12065    | Restart Process                                                 |
| 12070    | Job Running Succeeded                                           |
| 12071    | Job Running Failed                                              |
| 12072    | Job Killed                                                      |
| 12086    | Restart Agent                                                   |
| 12152    | Start Periodic Replication                                      |
| 12153    | Periodic Replication Completed                                  |
| 12154    | Start Streaming Replication                                     |
| 12155    | Restart Streaming Replication                                   |
| 12156    | Stop Streaming Replication                                      |
| 12157    | Skip Periodic Synchronization                                   |
| 14005    | NameNode Switchover                                             |
| 14028    | HDFS DiskBalancer Task                                          |
| 14029    | Active NameNode Entered Security Mode and Generated New Fsimage |

| Event ID | Event Name                                         |
|----------|----------------------------------------------------|
| 17001    | Oozie Workflow Execution Failure                   |
| 17002    | Oozie Scheduled Job Execution Failure              |
| 18001    | ResourceManager Switchover                         |
| 18004    | JobHistoryServer Switchover                        |
| 19001    | HMaster Failover                                   |
| 20003    | Hue Failover                                       |
| 24002    | Flume Channel Overflow                             |
| 25001    | LdapServer Failover                                |
| 27000    | DBServer Switchover                                |
| 29001    | Impala HaProxy Active/Standby Switchover           |
| 29002    | Impala StateStoreCatalog Active/Standby Switchover |
| 38003    | Adjust Topic Data Storage Period                   |
| 43014    | Spark2x Data Skew                                  |
| 43015    | Spark2x SQL Large Query Results                    |
| 43016    | Spark2x SQL Execution Timeout                      |
| 43024    | Start JDBCServer                                   |
| 43025    | Stop JDBCServer                                    |
| 43026    | ZooKeeper Connection Succeeded                     |
| 43027    | Zookeeper Connection Failed                        |
| 44003    | Coordinator Switchover                             |

## 7.11.2 Viewing Alarms of an MRS Cluster

You can view and clear alarms on MRS. Typically, the system automatically clears an alarm when the fault is rectified. If the fault has been rectified but the alarm is not automatically cleared, you can manually clear the alarm. You can view the latest 100,000 alarms (including uncleared, manually cleared, and automatically cleared alarms) on MRS. If the number of cleared alarms exceeds 100,000 and is about to reach 110,000, the system automatically dumps the earliest 10,000 cleared alarms to the dump path.

- For versions earlier than MRS 3.x, the path is **`${BIGDATA_HOME}/OMSV100R001C00x8664/workspace/data`** on the active management node.
- For MRS 3.x or later, the path is **`${BIGDATA_HOME}/om-server/OMS/workspace/data`** on the active management node.

A directory is automatically generated when alarms are dumped for the first time.

## Viewing and Clearing Alarms on the Management Console

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, select a running cluster and click its name to switch to the cluster details page.

**Step 3** Click **Alarms** and view the alarm information in the alarm list.

- The alarm list page displays the latest 10 alarms by default.
- You can filter all alarms of the same severity. The results include cleared and uncleared alarms.
- Click **Export All**. In the displayed **Export** dialog box, set **Save As** and click **OK**.

**Table 7-88** Alarm descriptions

| Parameter  | Description       |
|------------|-------------------|
| Alarm ID   | ID of an alarm.   |
| Alarm Name | Name of an alarm. |



| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Severity  | <p>Alarm severity.</p> <p>In versions earlier than MRS 3.x, the cluster alarm severity is as follows:</p> <ul style="list-style-type: none"> <li>● <b>Critical</b><br/>Indicates alarms reporting errors that affect cluster running, such as unavailable cluster services, node faults, data inconsistency between the active and standby GaussDB databases, and abnormal LdapServer data synchronization. You need to check the cluster status based on the alarms and rectify the faults in a timely manner.</li> <li>● <b>Major</b><br/>Indicates alarms reporting errors that affect some cluster functions, including process faults, periodic backup task failures, and abnormal key file permissions. Check the objects for which the alarms are generated based on the alarms and clear the alarms in a timely manner.</li> <li>● <b>Minor</b><br/>Indicates alarms reporting errors that do not affect major functions of the current cluster, including alarms indicating that the certificate file is about to expire, audit logs fail to be dumped, and the license file is about to expire.</li> <li>● <b>Warning</b><br/>Indicates an alarm of the lowest severity. It is used for information display or prompt and indicates that an event occurs in the scenarios when you stop a service, delete a service, stop an instance, delete an instance, delete a node, restart a service, restart an instance, perform an active/standby switchover for MRS Manager, scale in a host, or restore an instance. Additionally, this type of alarms also occurs when an instance is faulty, a job executed successfully, or a job failed to be executed.</li> </ul> <p>In MRS 3.x or later, the alarm severity of a cluster is as follows:</p> <ul style="list-style-type: none"> <li>● <b>Critical</b><br/>Indicates alarms reporting errors that affect cluster running, such as unavailable cluster services, node faults, data inconsistency between the active and standby GaussDB databases, and abnormal LdapServer data synchronization. You need to check the cluster status based on the alarms and rectify the faults in a timely manner.</li> <li>● <b>Major</b><br/>Indicates alarms reporting errors that affect some cluster functions, including process faults, periodic backup task failures, and abnormal key file permissions. Check the objects for which the alarms are generated based on the alarms and clear the alarms in a timely manner.</li> <li>● <b>Minor</b><br/>Indicates alarms reporting errors that do not affect major functions of the current cluster, including alarms indicating</li> </ul> |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | <p>that the certificate file is about to expire, audit logs fail to be dumped, and the license file is about to expire.</p> <ul style="list-style-type: none"> <li>Warning<br/>Indicates an alarm of the lowest severity. It is used for information display or prompt and indicates that an event occurs in the scenarios when you stop a service, delete a service, stop an instance, delete an instance, delete a node, restart a service, restart an instance, perform an active/standby switchover for MRS Manager, scale in a host, or restore an instance. Additionally, this type of alarms also occurs when an instance is faulty, a job executed successfully, or a job failed to be executed.</li> </ul> |
| Generated | Time when the alarm is generated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Location  | Details about the alarm.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Operation | <p>If the alarm can be manually cleared, click <b>Clear Alarm</b>.</p> <p>To view details about an alarm, click <b>View Help</b>. (This function is available in MRS 3.x or later).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Step 4** Click **Advanced Search**. In the displayed alarm search area, set search criteria and click **Search** to view information about specified alarms. Click **Reset** to clear the search criteria.

 **NOTE**

The start time and end time are specified in **Time Range**. You can search for alarms generated within the time range.

Handle the alarm by referring to **Alarm Reference**. If the alarms in some scenarios are generated due to other cloud services that MRS depends on, you need to contact maintenance personnel of the corresponding cloud services.

**Step 5** Click **Clear Alarm** if you need to. In the displayed dialog box, click **OK**.

 **NOTE**

After handling multiple alarms, you can select and clear one or more of them in batches by clicking **Clear Alarm**. Each batch can only clear a maximum of 300 alarms.


----End

## Viewing and Clearing Alarms on FusionInsight Manager (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Alarm > Alarms**.

**Step 3** View the alarm information reported by each cluster on FusionInsight Manager, including the alarm name, ID, severity, and generation time. By default, the latest 10 alarms are displayed on each page.

**Step 4** You can click  on the left of an alarm to view detailed alarm parameters. [Table 7-89](#) describes the parameters.

**Table 7-89** Alarm parameters

| Parameter              | Description                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Alarm ID               | Alarm ID.                                                                                                                                                                                                                                                                                                                                                                                         |
| Alarm Name             | Alarm name.                                                                                                                                                                                                                                                                                                                                                                                       |
| Alarm Severity         | Alarm severity. The options are <b>Critical</b> , <b>Major</b> , <b>Minor</b> , and <b>Suggestion</b> .                                                                                                                                                                                                                                                                                           |
| Generated              | Time when the alarm is generated.                                                                                                                                                                                                                                                                                                                                                                 |
| Cleared                | Time when an alarm is cleared. If the alarm is not cleared, -- is displayed.                                                                                                                                                                                                                                                                                                                      |
| Source                 | Cluster name.                                                                                                                                                                                                                                                                                                                                                                                     |
| Object                 | Service, process, or module that triggers the alarm.                                                                                                                                                                                                                                                                                                                                              |
| Auto Clear             | Whether the alarm can be automatically cleared after the fault is rectified.                                                                                                                                                                                                                                                                                                                      |
| Alarm Status           | Current status of the alarm. The options are <b>Auto</b> , <b>Manual</b> , and <b>Uncleared</b> .                                                                                                                                                                                                                                                                                                 |
| Alarm Cause            | Possible cause of an alarm.                                                                                                                                                                                                                                                                                                                                                                       |
| Serial Number          | Number of alarms generated by the system.                                                                                                                                                                                                                                                                                                                                                         |
| Additional Information | Error information.                                                                                                                                                                                                                                                                                                                                                                                |
| Location               | Detailed information for locating the alarm, which includes the following: <ul style="list-style-type: none"><li>● <b>Source</b>: cluster for which the alarm is generated.</li><li>● <b>ServiceName</b>: service for which the alarm is generated.</li><li>● <b>RoleName</b>: role for which the alarm is generated.</li><li>● <b>HostName</b>: host for which the alarm is generated.</li></ul> |

**Step 5** Manage alarms.

- Click **Export All** to export all alarm details.
- After handling multiple alarms, you can select and clear one or more of them in batches by clicking **Clear Alarm**. Each batch can only clear a maximum of 300 alarms.
- You can filter alarms by object or severity.
- You can click **Advanced Search** to search for alarms by alarm ID, name, type, start time, or end time. Click **Search** to filter alarms that meet the search

criteria. Click **Advanced Search** again to view the number of search criteria that you have configured.

- You can click **Clear**, **Mask**, or **View Help** to perform corresponding operations on an alarm.
- If there are a large number of alarms, you can click **View by Category** to sort uncleared alarms by alarm ID. After alarms are classified, click the number of uncleared alarms to view alarm details.

----End

## Viewing and Clearing Alarms on MRS Manager (MRS 2.x or Earlier)

**Step 1** On MRS Manager, click **Alarms** to view the alarm information in the alarm list.

- The alarm list page displays the latest 10 alarms by default.
- You can filter all alarms of the same severity in **Severity**. The results include cleared and uncleared alarms.

**Step 2** Click **Advanced Search**. In the displayed alarm search area, set search criteria and click **Search** to view information about specified alarms. Click **Reset** to clear the search criteria.

### NOTE

**Start Time** and **End Time** indicate the start time and end time of a time range. You can search for alarms generated within the time range.

Handle the alarm by referring to **Alarm Reference**. If the alarms in some scenarios are generated due to other cloud services that MRS depends on, you need to contact maintenance personnel of the corresponding cloud services.

**Step 3** Click **Clear Alarm** after the fault is rectified to manually clear the alarm.

### NOTE

After handling multiple alarms, you can select and clear one or more of them in batches by clicking **Clear Alarm**. Each batch can only clear a maximum of 300 alarms.

----End

## 7.11.3 Configuring Alarm Thresholds for an MRS Cluster

Manager allows you to configure monitoring metric thresholds to monitor the health of various metrics. If abnormal data occurs and meets the preset conditions, the system will trigger an alarm, which will appear on the alarm page.

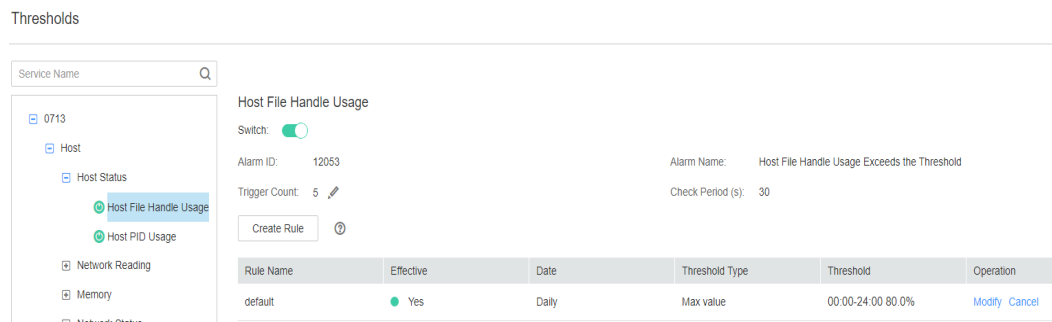
### Configuring Alarm Thresholds for an MRS Cluster (MRS 3.x or Later)

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Alarm > Thresholds**.

**Step 3** Select a monitoring metric for a host or service in the cluster.

**Figure 7-48** Configuring the threshold for a metric



For example, after selecting **Host Memory Usage**, the information about this indicator threshold is displayed.

- If the alarm sending switch is turned on, an alarm will be triggered if the threshold is reached.
- When **Alarm Severity** is on, hierarchical alarms are enabled. The system dynamically reports alarms at each severity based on the real-time metric values and hierarchical thresholds set for the severity. MRS 3.3.0 or later supports this function.
- **Alarm ID** and **Alarm Name**: alarm information triggered against the threshold
- **Trigger Count**: FusionInsight Manager checks whether the value of a monitoring metric reaches the threshold. If the number of consecutive checks reaches the value of **Trigger Count**, an alarm is generated. **Trigger Count** is configurable.
- **Check Period (s)**: interval for the system to check the monitoring metric.
- The rules in the rule list are used to trigger alarms.



**Step 4** Click **Create Rule** to add rules used for monitoring indicators.

**Table 7-90** Monitoring indicator rule parameters

| Parameter | Description               | Example Value                                                                                                     |
|-----------|---------------------------|-------------------------------------------------------------------------------------------------------------------|
| Rule Name | Name of a rule.           | CPU_MAX                                                                                                           |
| Severity  | Select an alarm severity. | <ul style="list-style-type: none"> <li>• Critical</li> <li>• Major</li> <li>• Minor</li> <li>• Warning</li> </ul> |

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                        | Example Value                                                                                 |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Threshold Type | You can use the maximum or minimum value of an indicator as the alarm triggering threshold. If <b>Threshold Type</b> is set to <b>Max value</b> , the system generates an alarm when the value of the specified indicator is greater than the threshold. If <b>Threshold Type</b> is set to <b>Min value</b> , the system generates an alarm when the value of the specified indicator is less than the threshold. | <ul style="list-style-type: none"> <li>• Max value</li> <li>• Min value</li> </ul>            |
| Date           | This parameter is used to set the date when the rule takes effect.                                                                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• Daily</li> <li>• Weekly</li> <li>• Others</li> </ul> |
| Add Date       | This parameter is available only when <b>Date</b> is set to <b>Others</b> . You can set the date when the rule takes effect. Multiple options are available.                                                                                                                                                                                                                                                       | 09-30                                                                                         |
| Thresholds     | This parameter is used to set the time range when the rule takes effect.                                                                                                                                                                                                                                                                                                                                           | Start and End Time:<br>00:00-08:30                                                            |
|                | Threshold of the rule monitoring metric                                                                                                                                                                                                                                                                                                                                                                            | Threshold: 10                                                                                 |

 **NOTE**

You can click  to set multiple time ranges for the threshold or click  to delete one.

**Step 5** Click **OK** to save the rules.

**Step 6** Locate the row that contains an added rule, and click **Apply** in the **Operation** column. The value of **Effective** for this rule changes to **Yes**.

A new rule can be applied only after you click **Cancel** for an existing rule.

----End

## Configuring Alarm Thresholds for an MRS Cluster (MRS 2.x or Earlier)

- Step 1** On MRS Manager, click **System**.
- Step 2** In **Configuration**, click **Configure Alarm Threshold** under **Monitoring and Alarm**, select monitoring metrics as planned, and set their baselines.
- Step 3** Click a metric, for example, **CPU Usage**, and click **Create Rule**.
- Step 4** In the displayed dialog box dialog box, set monitoring metric rule parameters.

**Table 7-91** Monitoring indicator rule parameters

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                        | Value                                                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| Rule Name      | Name of a rule.                                                                                                                                                                                                                                                                                                                                                                                                    | CPU_MAX                                                                                                      |
| Reference Date | Date on which the reference metric history is generated.                                                                                                                                                                                                                                                                                                                                                           | 2014/11/06<br>(example)                                                                                      |
| Threshold Type | You can use the maximum or minimum value of an indicator as the alarm triggering threshold. If <b>Threshold Type</b> is set to <b>Max value</b> , the system generates an alarm when the value of the specified indicator is greater than the threshold. If <b>Threshold Type</b> is set to <b>Min value</b> , the system generates an alarm when the value of the specified indicator is less than the threshold. | <ul style="list-style-type: none"><li>• Max value</li><li>• Min value</li></ul>                              |
| Severity       | Severity                                                                                                                                                                                                                                                                                                                                                                                                           | <ul style="list-style-type: none"><li>• Critical</li><li>• Major</li><li>• Minor</li><li>• Warning</li></ul> |
| Time Range     | Period in which the rule takes effect.                                                                                                                                                                                                                                                                                                                                                                             | From 00:00 to 23:59<br>(example)                                                                             |
| Threshold      | Threshold of the rule monitoring metric                                                                                                                                                                                                                                                                                                                                                                            | 80 (example)                                                                                                 |
| Date           | Type of date when the rule takes effect.                                                                                                                                                                                                                                                                                                                                                                           | <ul style="list-style-type: none"><li>• Workday</li><li>• Weekend</li><li>• Other</li></ul>                  |
| Add Date       | This parameter is valid only when <b>Date</b> is set to <b>Other</b> . You can select multiple dates.                                                                                                                                                                                                                                                                                                              | 11/30<br>(example)                                                                                           |

- Step 5** Click **OK**. A message is displayed in the upper right corner of the page, indicating that the template is saved successfully.

**Send alarm** is selected by default. **Trigger Count**: FusionInsight Manager checks whether the value of a monitoring metric reaches the threshold. If the number of

consecutive checks reaches the value of **Trigger Count**, an alarm is generated. **Trigger Count** is configurable. **Check Period (s)** indicates the interval at which MRS Manager checks monitoring metrics.

- Step 6** Locate the row that contains the newly added rule and click **Apply** in the **Operation** column. A message is displayed in the upper right corner, indicating that the rule *xx* is successfully added. Click **Cancel** in the **Operation** column. A message is displayed in the upper right corner, indicating that the rule *xx* is successfully canceled.

----End

## Monitoring Metric Reference (MRS 3.x or Later)

FusionInsight Manager alarm monitoring metrics are classified as node information metrics and cluster service metrics. [Table 7-92](#) lists the metrics whose thresholds can be configured for a node, and [Table 7-93](#) lists the metrics whose thresholds can be configured for a component.

**Table 7-92** Node monitoring metrics

| Metric Group | Metric            | Description                                                                                                                                                                                                           | Default Threshold |
|--------------|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| CPU          | Host CPU Usage    | This indicator reflects the computing and control capabilities of the current cluster in a measurement period. By observing the indicator value, you can better understand the overall resource usage of the cluster. | 90.0%             |
| Disk         | Disk Usage        | Indicates the disk usage of a host.                                                                                                                                                                                   | 90.0%             |
|              | Disk Inode Usage  | Indicates the disk inode usage in a measurement period.                                                                                                                                                               | 80.0%             |
| Memory       | Host Memory Usage | Indicates the average memory usage at the current time.                                                                                                                                                               | 90.0%             |



| Metric Group    | Metric                   | Description                                                                                            | Default Threshold |
|-----------------|--------------------------|--------------------------------------------------------------------------------------------------------|-------------------|
| Host Status     | Host File Handle Usage   | Indicates the usage of file handles of the host in a measurement period.                               | 80.0%             |
|                 | Host PID Usage           | Indicates the PID usage of a host.                                                                     | 90%               |
| Network Status  | TCP Ephemeral Port Usage | Indicates the usage of temporary TCP ports of the host in a measurement period.                        | 80.0%             |
| Network Reading | Read Packet Error Rate   | Indicates the read packet error rate of the network interface on the host in a measurement period.     | 0.5%              |
|                 | Read Packet Dropped Rate | Indicates the read packet dropped rate of the network interface on the host in a measurement period.   | 0.5%              |
|                 | Read Throughput Rate     | Indicates the average read throughput (at MAC layer) of the network interface in a measurement period. | 80%               |
| Network Writing | Write Packet Error Rate  | Indicates the write packet error rate of the network interface on the host in a measurement period.    | 0.5%              |

| Metric Group | Metric                        | Description                                                                                             | Default Threshold |
|--------------|-------------------------------|---------------------------------------------------------------------------------------------------------|-------------------|
|              | Write Packet Dropped Rate     | Indicates the write packet dropped rate of the network interface on the host in a measurement period.   | 0.5%              |
|              | Write Throughput Rate         | Indicates the average write throughput (at MAC layer) of the network interface in a measurement period. | 80%               |
| Process      | Uninterruptible Sleep Process | Number of D state processes on the host in a measurement period                                         | 0                 |
|              | omm Process Usage             | omm process usage in a measurement period                                                               | 90                |

**Table 7-93** Cluster service indicators

| Service   | Metric Group | Metric Name                                 | Metric Description                     | Default Threshold |
|-----------|--------------|---------------------------------------------|----------------------------------------|-------------------|
| DBService | Database     | Usage of the Number of Database Connections | Usage of database connections          | 90%               |
|           |              | Disk Space Usage of the Data Directory      | Disk space usage of the data directory | 80%               |
| Flume     | Agent        | Heap Memory Usage Calculate                 | Flume heap memory usage                | 95.0%             |

| Service | Metric Group   | Metric Name                                 | Metric Description                  | Default Threshold |
|---------|----------------|---------------------------------------------|-------------------------------------|-------------------|
|         |                | Flume Direct Memory Usage Statistics        | Flume direct memory usage           | 80.0%             |
|         |                | Flume Non-heap Memory Usage                 | Flume non-heap memory usage         | 80.0%             |
|         |                | Total GC duration of Flume process          | Flume total GC time                 | 12000 ms          |
| HBase   | GC             | GC time for old generation                  | Total GC time of RegionServer       | 5000 ms           |
|         |                | GC time for old generation                  | Total GC time of HMaster            | 5000 ms           |
|         | CPU & memory   | RegionServer Direct Memory Usage Statistics | RegionServer direct memory usage    | 90%               |
|         |                | RegionServer Heap Memory Usage Statistics   | RegionServer heap memory usage      | 90%               |
|         |                | HMaster Direct Memory Usage                 | HMaster direct memory usage         | 90%               |
|         |                | HMaster Heap Memory Usage Statistics        | HMaster heap memory usage           | 90%               |
|         | <b>Service</b> | Number of Online Regions of a RegionServer  | Number of regions of a RegionServer | 2000              |

| Service | Metric Group   | Metric Name                                                  | Metric Description                                                           | Default Threshold |
|---------|----------------|--------------------------------------------------------------|------------------------------------------------------------------------------|-------------------|
|         |                | Region in transaction count over threshold                   | Number of regions that are in the RIT state and reach the threshold duration | 1                 |
|         | Replication    | Replication sync failed times (RegionServer )                | Number of times that DR data fails to be synchronized                        | 1                 |
|         |                | Number of Log Files to Be Synchronized in the Active Cluster | Number of log files to be synchronized in the active cluster                 | 128               |
|         |                | Number of HFiles to Be Synchronized in the Active Cluster    | Number of HFiles to be synchronized in the active cluster                    | 128               |
|         | Queue          | Compaction Queue Size                                        | Size of the Compaction queue                                                 | 100               |
| HDFS    | File and Block | Lost Blocks                                                  | Number of block copies that the HDFS lacks of                                | 0                 |
|         |                | Blocks Under Replicated                                      | Total number of blocks that need to be replicated by the NameNode            | 1000              |
|         | RPC            | Average Time of Active NameNode RPC Processing               | Average NameNode RPC processing time                                         | 100 ms            |

| Service | Metric Group | Metric Name                                               | Metric Description                                                                                   | Default Threshold |
|---------|--------------|-----------------------------------------------------------|------------------------------------------------------------------------------------------------------|-------------------|
|         |              | Average Time of Active NameNode RPC Queuing               | Average NameNode RPC queuing time                                                                    | 200 ms            |
|         | Disk         | HDFS Disk Usage                                           | HDFS disk usage                                                                                      | 80%               |
|         |              | DataNode Disk Usage                                       | Disk usage of DataNodes in the HDFS                                                                  | 80%               |
|         |              | Percentage of Reserved Space for Replicas of Unused Space | Percentage of the reserved disk space of all the copies to the total unused disk space of DataNodes. | 90%               |
|         | Resource     | Faulty DataNodes                                          | Indicates the number of faulty DataNodes.                                                            | 3                 |
|         |              | NameNode Non-Heap Memory Usage Statistics                 | Indicates the percentage of NameNode non-heap memory usage.                                          | 90%               |
|         |              | NameNode Direct Memory Usage Statistics                   | Indicates the percentage of direct memory used by NameNodes.                                         | 90%               |
|         |              | NameNode Heap Memory Usage Statistics                     | Indicates the percentage of NameNode non-heap memory usage.                                          | 95%               |
|         |              | DataNode Direct Memory Usage Statistics                   | Indicates the percentage of direct memory used by DataNodes.                                         | 90%               |

| Service | Metric Group       | Metric Name                                                         | Metric Description                                                                 | Default Threshold |
|---------|--------------------|---------------------------------------------------------------------|------------------------------------------------------------------------------------|-------------------|
|         |                    | DataNode Heap Memory Usage Statistics                               | DataNode heap memory usage                                                         | 95%               |
|         |                    | DataNode Heap Memory Usage Statistics                               | Indicates the percentage of DataNode non-heap memory usage.                        | 90%               |
|         | Garbage Collection | GC Time (NameNode)/ GC Time (DataNode)                              | Indicates the Garbage collection (GC) duration of NameNodes per minute.            | 12000 ms          |
|         |                    | GC Time                                                             | Indicates the GC duration of DataNodes per minute.                                 | 12000 ms          |
| Hive    | HQL                | Percentage of HQL Statements That Are Executed Successfully by Hive | Indicates the percentage of HQL statements that are executed successfully by Hive. | 90.0%             |
|         | Background         | Background Thread Usage                                             | Background thread usage                                                            | 90%               |
|         | GC                 | Total GC time of MetaStore                                          | Indicates the total GC time of MetaStore.                                          | 12000 ms          |
|         |                    | Total GC Time in Milliseconds                                       | Indicates the total GC time of HiveServer.                                         | 12000 ms          |
|         | Capacity           | Percentage of HDFS Space Used by Hive to the Available Space        | Indicates the percentage of HDFS space used by Hive to the available space.        | 85.0%             |

| Service | Metric Group | Metric Name                                                                                                | Metric Description                                                                                                                          | Default Threshold |
|---------|--------------|------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
|         | CPU & memory | MetaStore Direct Memory Usage Statistics                                                                   | MetaStore direct memory usage                                                                                                               | 95%               |
|         |              | MetaStore Non-Heap Memory Usage Statistics                                                                 | MetaStore non-heap memory usage                                                                                                             | 95%               |
|         |              | MetaStore Heap Memory Usage Statistics                                                                     | MetaStore heap memory usage                                                                                                                 | 95%               |
|         |              | HiveServer Direct Memory Usage Statistics                                                                  | HiveServer direct memory usage                                                                                                              | 95%               |
|         |              | HiveServer Non-Heap Memory Usage Statistics                                                                | HiveServer non-heap memory usage                                                                                                            | 95%               |
|         |              | HiveServer Heap Memory Usage Statistics                                                                    | HiveServer heap memory usage                                                                                                                | 95%               |
|         | Session      | Percentage of Sessions Connected to the HiveServer to Maximum Number of Sessions Allowed by the HiveServer | Indicates the percentage of the number of sessions connected to the HiveServer to the maximum number of sessions allowed by the HiveServer. | 90.0%             |

| Service | Metric Group | Metric Name                                                   | Metric Description                                                                               | Default Threshold |
|---------|--------------|---------------------------------------------------------------|--------------------------------------------------------------------------------------------------|-------------------|
| Kafka   | Partition    | Percentage of Partitions That Are Not Completely Synchronized | Indicates the percentage of partitions that are not completely synchronized to total partitions. | 50%               |
|         | Others       | Unavailable Partition Percentage                              | Percentage of unavailable partitions of each Kafka topic                                         | 40%               |
|         |              | User Connection Usage on Broker                               | Usage of user connections on Broker                                                              | 80%               |
|         | Disk         | Broker Disk Usage                                             | Indicates the disk usage of the disk where the Broker data directory is located.                 | 80.0%             |
|         |              | Disk I/O Rate of a Broker                                     | I/O usage of the disk where the Broker data directory is located                                 | 80%               |
|         | Process      | Broker GC Duration per Minute                                 | Indicates the GC duration of the Broker process per minute.                                      | 12000 ms          |
|         |              | Heap Memory Usage of Kafka                                    | Indicates the Kafka heap memory usage.                                                           | 95%               |
|         |              | Kafka Direct Memory Usage                                     | Indicates the Kafka direct memory usage.                                                         | 95%               |



| Service   | Metric Group       | Metric Name                                       | Metric Description                                    | Default Threshold |
|-----------|--------------------|---------------------------------------------------|-------------------------------------------------------|-------------------|
| Loader    | Memory             | Heap Memory Usage Calculate                       | Indicates the Loader heap memory usage.               | 95%               |
|           |                    | Direct Memory Usage of Loader                     | Indicates the Loader direct memory usage.             | 80.0%             |
|           |                    | Non-heap Memory Usage of Loader                   | Indicates the Loader non-heap memory usage.           | 80%               |
|           | GC                 | Total GC time of Loader                           | Indicates the total GC time of Loader.                | 12000 ms          |
| MapReduce | Garbage Collection | GC Time                                           | Indicates the GC time.                                | 12000 ms          |
|           | Resource           | JobHistoryServer Direct Memory Usage Statistics   | Indicates the JobHistoryServer direct memory usage.   | 90%               |
|           |                    | JobHistoryServer Non-Heap Memory Usage Statistics | Indicates the JobHistoryServer non-heap memory usage. | 90%               |
|           |                    | JobHistoryServer Heap Memory Usage Statistics     | Indicates the JobHistoryServer non-heap memory usage. | 95%               |
| Oozie     | Memory             | Oozie Heap Memory Usage Calculate                 | Indicates the Oozie heap memory usage.                | 95.0%             |
|           |                    | Oozie Direct Memory Usage                         | Indicates the Oozie direct memory usage.              | 80.0%             |

| Service           | Metric Group | Metric Name                                   | Metric Description                         | Default Threshold |
|-------------------|--------------|-----------------------------------------------|--------------------------------------------|-------------------|
|                   |              | Oozie Non-heap Memory Usage                   | Indicates the Oozie non-heap memory usage. | 80%               |
|                   | GC           | Total GC duration of Oozie                    | Indicates the Oozie total GC time.         | 12000 ms          |
| Spark/<br>Spark2x | Memory       | JDBCServer2x Heap Memory Usage Statistics     | JDBCServer2x heap memory usage             | 95%               |
|                   |              | JDBCServer2x Direct Memory Usage Statistics   | JDBCServer2x direct memory usage           | 95%               |
|                   |              | JDBCServer2x Non-Heap Memory Usage Statistics | JDBCServer2x non-heap memory usage         | 95%               |
|                   |              | JobHistory2x Direct Memory Usage Statistics   | JobHistory2x direct memory usage           | 95%               |
|                   |              | JobHistory2x Non-Heap Memory Usage Statistics | JobHistory2x non-heap memory usage         | 95%               |
|                   |              | JobHistory2x Heap Memory Usage Statistics     | JobHistory2x heap memory usage             | 95%               |
|                   |              | IndexServer2x Direct Memory Usage Statistics  | IndexServer2x direct memory usage          | 95%               |

| Service | Metric Group | Metric Name                                    | Metric Description                  | Default Threshold                                                                              |       |
|---------|--------------|------------------------------------------------|-------------------------------------|------------------------------------------------------------------------------------------------|-------|
|         |              | IndexServer2x Heap Memory Usage Statistics     | IndexServer2x heap memory usage     | 95%                                                                                            |       |
|         |              | IndexServer2x Non-Heap Memory Usage Statistics | IndexServer2x non-heap memory usage | 95%                                                                                            |       |
|         | GC Count     | Full GC Number of JDBCServer2x                 | Full GC times of JDBCServer2x       | 12                                                                                             |       |
|         |              | Full GC Number of JobHistory2x                 | Full GC times of JobHistory2x       | 12                                                                                             |       |
|         |              | Full GC Number of IndexServer2x                | Full GC times of IndexServer2x      | 12                                                                                             |       |
|         | GC Time      | Total GC Time in Milliseconds                  | Total GC time of JDBCServer2x       | 12000 ms                                                                                       |       |
|         |              | Total GC Time in Milliseconds                  | Total GC time of JobHistory2x       | 12000 ms                                                                                       |       |
|         |              | Total GC Time in Milliseconds                  | Total GC time of IndexServer2x      | 12000 ms                                                                                       |       |
|         | Storm        | Cluster                                        | Number of Available Supervisors     | Indicates the number of available Supervisor processes in the cluster in a measurement period. | 1     |
|         |              |                                                | Slot Usage                          | Indicates the slot usage in the cluster in a measurement period.                               | 80.0% |

| Service | Metric Group       | Metric Name                                      | Metric Description                                              | Default Threshold |
|---------|--------------------|--------------------------------------------------|-----------------------------------------------------------------|-------------------|
|         | Nimbus             | Nimbus Heap Memory Usage Calculate               | Indicates the Nimbus heap memory usage.                         | 80%               |
| Yarn    | Resources          | NodeManager Direct Memory Usage Statistics       | Indicates the percentage of direct memory used by NodeManagers. | 90%               |
|         |                    | NodeManager Heap Memory Usage Statistics         | Indicates the percentage of NodeManager heap memory usage.      | 95%               |
|         |                    | NodeManager Non-Heap Memory Usage Statistics     | Indicates the percentage of NodeManager non-heap memory usage.  | 90%               |
|         |                    | ResourceManager Direct Memory Usage Statistics   | Indicates the Kafka direct memory usage.                        | 90%               |
|         |                    | ResourceManager Heap Memory Usage Statistics     | Indicates the ResourceManager heap memory usage.                | 95%               |
|         |                    | ResourceManager Non-Heap Memory Usage Statistics | Indicates the ResourceManager non-heap memory usage.            | 90%               |
|         | Garbage collection | GC Time                                          | Indicates the GC duration of NodeManager per minute.            | 12000 ms          |

| Service   | Metric Group | Metric Name                           | Metric Description                                                                      | Default Threshold |
|-----------|--------------|---------------------------------------|-----------------------------------------------------------------------------------------|-------------------|
|           |              | GC Time                               | Indicates the GC duration of ResourceManager per minute.                                | 12000 ms          |
|           | Others       | Failed Applications of root queue     | Number of failed tasks in the root queue                                                | 50                |
|           |              | Terminated Applications of root queue | Number of killed tasks in the root queue                                                | 50                |
|           | CPU & memory | Pending Memory                        | Pending memory capacity                                                                 | 83886080 MB       |
|           | Application  | Pending Applications                  | Pending tasks                                                                           | 60                |
| ZooKeeper | Connection   | ZooKeeper Connections Usage           | Indicates the percentage of the used connections to the total connections of ZooKeeper. | 80%               |
|           | CPU & memory | Directmemory Usage Calculate          | Indicates the ZooKeeper heap memory usage.                                              | 95%               |
|           |              | Heap Memory Usage Calculate           | Indicates the ZooKeeper direct memory usage.                                            | 80%               |
|           | GC           | ZooKeeper GC Duration per Minute      | Indicates the GC time of ZooKeeper every minute.                                        | 12000 ms          |
| Ranger    | GC           | UserSync GC Duration                  | UserSync garbage collection (GC) duration                                               | 12000 ms          |
|           |              | RangerAdmin GC Duration               | RangerAdmin GC duration                                                                 | 12000 ms          |

| Service    | Metric Group | Metric Name                       | Metric Description                                   | Default Threshold                                         |
|------------|--------------|-----------------------------------|------------------------------------------------------|-----------------------------------------------------------|
|            |              | TagSync GC Duration               | TagSync GC duration                                  | 12000 ms                                                  |
|            | CPU & memory | UserSync Non-Heap Memory Usage    | UserSync non-heap memory usage                       | 80.0%                                                     |
|            |              | UserSync Direct Memory Usage      | UserSync direct memory usage                         | 80.0%                                                     |
|            |              | UserSync Heap Memory Usage        | UserSync heap memory usage                           | 95.0%                                                     |
|            |              | RangerAdmin Non-Heap Memory Usage | RangerAdmin non-heap memory usage                    | 80.0%                                                     |
|            |              | RangerAdmin Heap Memory Usage     | RangerAdmin heap memory usage                        | 95.0%                                                     |
|            |              | RangerAdmin Direct Memory Usage   | RangerAdmin direct memory usage                      | 80.0%                                                     |
|            |              | TagSync Direct Memory Usage       | TagSync direct memory usage                          | 80.0%                                                     |
|            |              | TagSync Non-Heap Memory Usage     | TagSync non-heap memory usage                        | 80.0%                                                     |
|            |              | TagSync Heap Memory Usage         | TagSync heap memory usage                            | 95.0%                                                     |
| ClickHouse |              | Cluster Quota                     | Clickhouse service quantity quota usage in ZooKeeper | Quota of the ZooKeeper nodes used by a ClickHouse service |

| Service | Metric Group | Metric Name                                                 | Metric Description                                                   | Default Threshold |
|---------|--------------|-------------------------------------------------------------|----------------------------------------------------------------------|-------------------|
|         |              | Capacity quota usage of the Clickhouse service in ZooKeeper | Capacity quota of ZooKeeper directory used by the ClickHouse service | 90%               |
| IoTDB   | GC           | IoTDBServer GC Duration                                     | IoTDBServer garbage collection (GC) duration                         | 12000 ms          |
|         | CPU & memory | IoTDBServer Heap Memory Usage                               | IoTDBServer heap memory usage                                        | 90%               |
|         |              | IoTDBServer Direct Memory Usage                             | IoTDBServer direct memory usage                                      | 90%               |

## 7.11.4 Configuring Alarm Masking for an MRS Cluster

If you do not want FusionInsight Manager to report specified alarms in the following scenarios, you can manually mask the alarms.

- Some unimportant alarms and minor alarms need to be masked.
- When a third-party product is integrated with MRS, some alarms of the product are duplicated with the alarms of MRS and need to be masked.
- When the deployment environment is special, certain alarms may be falsely reported and need to be masked.

Once an alarm is masked, any new alarms with the same ID will no longer appear on the **Alarm Management** page or be counted. However, previously reported alarms will still be displayed.

### NOTE

This section applies only to MRS 3.x or later.

**Step 1** Log in to FusionInsight Manager.

**Step 2** Choose **O&M > Alarm > Masking Setting**.

**Step 3** In the **Masking Setting** area, select the specified service or module.

**Step 4** Select an alarm from the alarm list.

**Figure 7-49** Masking an alarm

The information about the alarm is displayed, including the alarm name, ID, severity, masking status, and operations can be performed on the alarm.

- The masking status includes **Display** and **Masking**.
- Operations include **Masking** and **Help**.

#### NOTE

You can filter specified alarms based on the masking status and alarm severity.

**Step 5** Set the masking status for an alarm:

- Click **Masking**. In the displayed dialog box, click **OK** to change the alarm masking status to **Masking**.
- Click **Cancel Masking**. In the dialog box that is displayed, click **OK** to change the masking status of the alarm to **Display**.

----End

## 7.11.5 Connecting an MRS Cluster to SNMP to Report Alarms

If users need to view alarms and monitoring data of a cluster on the O&M platform, you can use Simple Network Management Protocol (SNMP) on FusionInsight Manager to report related data to the network management system (NMS).

### Prerequisites

The ECS corresponding to the server must be in the same VPC as the MRS cluster's Master node, and the Master node must be able to access the server's IP address and specified port.

### Connecting an MRS Cluster to SNMP to Report Alarms

**Step 1** Log in to the Manager and go to the SNMP configuration page.

- MRS 3.x or later:
  - a. Log in to FusionInsight Manager.
  - b. Choose **System > Interconnection > SNMP**.
  - c. Enable the SNMP service.
- Versions earlier than MRS 3.x:
  - a. Log in to MRS Manager and click **System**.
  - b. In **Configuration**, enable **Configure SNMP** under **Monitoring and Alarm**.



**Step 2** Set interconnection parameters based on [Table 7-94](#).**Table 7-94** Interconnection parameters

| Parameter               | Description                                                                                                                                                                                                                                                    |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Version                 | Specifies the version of SNMP, which can be: <ul style="list-style-type: none"><li>• <b>V2C</b>: This is an earlier version with low security.</li><li>• <b>V3</b>: This is a later version with higher security than SNMPv2c</li></ul> SNMPv3 is recommended. |
| Local Port              | Specifies the local port. The default value is <b>20000</b> . The value ranges from <b>1025</b> to <b>65535</b> .                                                                                                                                              |
| Read Community Name     | Specifies the read-only community name. This parameter is available only when <b>Version</b> is set to <b>V2C</b> .                                                                                                                                            |
| Write Community Name    | Specifies the write community name. This parameter is available only when <b>Version</b> is set to <b>V2C</b> .                                                                                                                                                |
| Security Username       | Specifies the SNMP security username. This parameter is available only when <b>Version</b> is set to <b>V3</b> .                                                                                                                                               |
| Authentication Protocol | Specifies the authentication protocol. This parameter is available only when <b>Version</b> is set to <b>V3</b> . SHA is recommended.                                                                                                                          |
| Authentication Password | Specifies the authentication password. This parameter is available only when <b>Version</b> is set to <b>V3</b> .                                                                                                                                              |
| Confirm Password        | Used to confirm the authentication password. This parameter is available only when <b>Version</b> is set to <b>V3</b> .                                                                                                                                        |
| Encryption Protocol     | Specifies the encryption protocol. This parameter is available only when <b>Version</b> is set to <b>V3</b> . AES256 is recommended.                                                                                                                           |
| Encryption Password     | Specifies the encryption password. This parameter is available only when <b>Version</b> is set to <b>V3</b> .                                                                                                                                                  |
| Confirm Password        | Used to confirm the encryption password. This parameter is available only when <b>Version</b> is set to <b>V3</b> .                                                                                                                                            |

 NOTE

- The value of **Security Username** cannot contain repeated strings with the unit length as a common factor of 64 (such as 1, 2, 4, and 8), for example, **abab** and **abcdabcd**.
- The **Authentication Password** and **Encryption Password** must contain 8 to 16 characters, including at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters. The two passwords must be different. The two passwords cannot be the same as the security username or the reverse of the security username.
- For security purposes, periodically change the authentication password and encryption password when the SNMP protocol is used.
- If SNMP v3 is used, a security user will be locked after five consecutive authentication failures within 5 minutes. The user will be automatically unlocked 5 minutes later.

**Step 3** Click **Create Trap Target** in the **Trap Target** area. In the displayed dialog box, set the following parameters:

- **Target Symbol:** specifies the trap target ID, which is the ID of the NMS or host that receives traps. The value consists of 1 to 255 characters, including letters or digits.
- **Target IP Address Mode** (only for MRS 3.x or later): mode of the target IP address. The options are **IPv4** and **IPv6**.
- **Target IP Address:** specifies the target IP address, which can communicate with the management plane IP address of the management node.
- **Target Port:** specifies the port receiving traps. The port number must be consistent with the peer end and ranges from 0 to 65535.
- **Trap Community Name:** This parameter is available only when **Version** is set to **V2C** and is used to report the community name.

Click **OK**.

The **Create Trap Target** dialog box is closed.

**Step 4** Click **OK**.

----End

## 7.11.6 Connecting an MRS Cluster to the Syslog Server to Report Alarms

If users need to view alarms and events of a cluster on the unified alarm reporting platform, you can use the Syslog protocol on FusionInsight Manager to report related data to the alarm platform.

---

**NOTICE**

If the Syslog protocol is not encrypted, data may be stolen.

---

### Prerequisites

The ECS corresponding to the server must be in the same VPC as the MRS cluster's Master node, and the Master node must be able to access the server's IP address and specified port.

## Connecting an MRS Cluster to the Syslog Server to Report Alarms

**Step 1** Log in to the Manager and go to the SNMP configuration page.

- MRS 3.x or later:
  - a. Log in to FusionInsight Manager.
  - b. Choose **System > Interconnection > Syslog**.
  - c. Enable the Syslog service.
- Versions earlier than MRS 3.x:
  - a. Log in to MRS Manager and click **System**.
  - b. In **Configuration**, enable **Configure Syslog** under **Monitoring and Alarm**.

**Step 2** Set northbound parameters based on [Table 7-95](#).

**Table 7-95** Syslog interconnection parameters

| Parameter Area  | Parameter              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syslog Protocol | Server IP Address Mode | Specifies the IP address mode of the interconnected server. The value can be <b>IPV4</b> or <b>IPV6</b> . (only for MRS 3.x or later)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                 | Server IP Address      | Specifies the IP address of the interconnected server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                 | Server Port            | Specifies the port number for interconnection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                 | Protocol               | Specifies the protocol type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>TCP</b></li> <li>• <b>UDP</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|                 | Severity Level         | Specifies the severity of the reported message. The options are as follows: <ul style="list-style-type: none"> <li>• <b>Emergency</b></li> <li>• <b>Alert</b></li> <li>• <b>Critical</b></li> <li>• <b>Error</b></li> <li>• <b>Warning</b></li> <li>• <b>Notice</b></li> <li>• <b>Informational</b> (default value)</li> <li>• <b>Debug</b></li> </ul> <p><b>NOTE</b><br/> <b>Severity Level</b> and <b>Facility</b> determine the priority of the sent message.</p> <p><b>Priority = Facility × 8 + Severity Level</b></p> <p>For the values of <b>Severity Level</b> and <b>Facility</b>, see <a href="#">Table 7-96</a>.</p> |

| Parameter Area            | Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                           | Facility                           | Specifies the module where the log is generated. For the available values of this parameter, see <a href="#">Table 7-96</a> . Default value <b>local use 0 (local0)</b> is recommended.                                                                                                                                                                                                                                                                                                                                                                                                                      |
|                           | Identifier                         | <p>Product ID.</p> <p>The identifier can contain a maximum of 256 characters, including letters, digits, underscores (_), periods (.), hyphens (-), spaces, and the following special characters:   \$ { }</p> <ul style="list-style-type: none"> <li>• MRS 3.x or later: The default value is <b>FusionInsight Manager</b>.</li> <li>• Versions earlier than MRS 3.x: The default value is <b>MRS Manager</b>.</li> </ul>                                                                                                                                                                                   |
| Report Message            | Report Format                      | <p>Specifies the message format of the alarm report. For details, see the help information on the page.</p> <p>The report format can contain a maximum of 1024 characters, including letters, digits, underscores (_), periods (.), hyphens (-), spaces, and the following special characters:   \$ { }</p> <p><b>NOTE</b><br/>For details about each field in the report format, see <a href="#">Table 7-97</a>.</p>                                                                                                                                                                                        |
|                           | Alarm Type                         | <p>Specifies the type of the alarm to be reported.</p> <ul style="list-style-type: none"> <li>• For MRS 3.x or later, see <a href="#">Table 7-97</a>.</li> <li>• Versions earlier than MRS 3.x: <ul style="list-style-type: none"> <li>- <b>Fault</b>: indicates that the Syslog alarm message is reported when MRS Manager generates an alarm.</li> <li>- <b>Clear</b>: indicates that a Syslog alarm message is reported when an alarm on MRS Manager is cleared.</li> <li>- <b>Event</b>: indicates that the Syslog alarm message is reported when MRS Manager generates an event.</li> </ul> </li> </ul> |
|                           | Alarm Severities                   | <p>Specifies the level of the alarm to be reported.</p> <ul style="list-style-type: none"> <li>• For MRS 3.x or later, see <a href="#">Table 7-97</a>.</li> <li>• Versions earlier than MRS 3.x: <b>Suggestion, Minor, Major, and Critical</b>.</li> </ul>                                                                                                                                                                                                                                                                                                                                                   |
| Uncleared Alarm Reporting | Periodic Uncleared Alarm Reporting | Specifies whether to report uncleared alarms in a specified period. You can toggle on or off the function. The function is toggled off by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Parameter Area     | Parameter                    | Description                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | Report Interval (min)        | Specifies the interval for periodically reporting uncleared alarms. This parameter is valid only when <b>Periodic Uncleared Alarm Reporting</b> is enabled. The default value is <b>15</b> , in minutes. The value ranges from <b>5</b> to <b>1440</b> (one day).                                                  |
| Heartbeat Settings | Heartbeat Reporting          | Specifies whether to periodically report Syslog heartbeat messages. You can toggle on or off the function. The function is toggled off by default.                                                                                                                                                                 |
|                    | Heartbeat Interval (minutes) | Specifies the interval for periodically reporting heartbeat messages. This parameter is valid only when <b>Heartbeat Reporting</b> is enabled. The default value is <b>15</b> , in minutes. The value ranges from <b>1</b> to <b>60</b> .                                                                          |
|                    | Heartbeat Packet             | Specifies the heartbeat message to be reported. This parameter is valid when <b>Heartbeat Reporting</b> is toggled on and cannot be left blank. The value can contain a maximum of 256 characters, including digits, letters, underscores (_), vertical bars ( ), colons (:), spaces, commas (,), and periods (.). |

 **NOTE**

After the periodic heartbeat packet function is enabled, packets may be interrupted during automatic recovery of some cluster error tolerance (for example, active/standby OMS switchover). In this case, wait for automatic recovery.

**Step 3** Click **OK**.

----End

## Related Information

**Table 7-96** Numeric codes of **Severity Level** and **Facility**

| Severity Level   | Facility                                 | Numeric Code |
|------------------|------------------------------------------|--------------|
| <b>Emergency</b> | kernel messages                          | 0            |
| <b>Alert</b>     | user-level messages                      | 1            |
| <b>Critical</b>  | mail system                              | 2            |
| <b>Error</b>     | system daemons                           | 3            |
| <b>Warning</b>   | security/authorization messages (note 1) | 4            |
| <b>Notice</b>    | messages generated internally by syslog  | 5            |

| Severity Level | Facility                                 | Numeric Code |
|----------------|------------------------------------------|--------------|
| Informational  | line printer subsystem                   | 6            |
| <b>Debug</b>   | network news subsystem                   | 7            |
| -              | UUCP subsystem                           | 8            |
| -              | clock daemon (note 2)                    | 9            |
| -              | security/authorization messages (note 1) | 10           |
| -              | FTP daemon                               | 11           |
| -              | NTP subsystem                            | 12           |
| -              | log audit (note 1)                       | 13           |
| -              | log alert (note 1)                       | 14           |
| -              | clock daemon (note 2)                    | 15           |
| -              | local use 0~7 (local0 ~ local7)          | 16 to 23     |

**Table 7-97** Report format information fields

| Information Field | Description                                                                                                                                              |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| dn                | Cluster name                                                                                                                                             |
| id                | Alarm ID                                                                                                                                                 |
| name              | Alam name                                                                                                                                                |
| serialNo          | Alarm serial number<br><b>NOTE</b><br>The serial numbers of the fault alarms and the corresponding clear alarms are the same.                            |
| category          | Alarm type. The options are as follows: <ul style="list-style-type: none"> <li>● 0: fault alarm</li> <li>● 1: clear alarm</li> <li>● 2: event</li> </ul> |
| occurTime         | Time when the alarm was generated                                                                                                                        |
| clearTime         | Time when this alarm was cleared                                                                                                                         |
| isAutoClear       | Whether an alarm is automatically cleared. The options are as follows: <ul style="list-style-type: none"> <li>● 1: yes</li> <li>● 0: no</li> </ul>       |
| locationInfo      | Location where the alarm was generated                                                                                                                   |

| Information Field | Description                                                                                                                                                                            |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| clearType         | Alarm clearance type. The options are as follows: <ul style="list-style-type: none"><li>• -1: not cleared</li><li>• 0: automatically cleared</li><li>• 2: manually cleared</li></ul>   |
| level             | Severity. The options are as follows: <ul style="list-style-type: none"><li>• 1: critical alarm</li><li>• 2: major alarm</li><li>• 3: minor alarm</li><li>• 4: warning alarm</li></ul> |
| cause             | Alarm cause                                                                                                                                                                            |
| additionalInfo    | Additional information                                                                                                                                                                 |
| object            | Alarm object                                                                                                                                                                           |

## 7.11.7 Periodically Backing Up Alarm and Audit Information

You can modify the configuration file to periodically back up FusionInsight Manager alarm information, FusionInsight Manager audit information, and audit information of all services to the specified storage location.

The backup can be performed using FTP or SFTP. FTP does not encrypt data, which may cause security risks. Therefore, SFTP is recommended.

### NOTE

This section applies only to MRS 3.x or later.

**Step 1** Log in to the active management node as user **omm**.

### NOTE

Perform this operation only on the active management node. Scheduled backup is not supported on the standby management node.

**Step 2** Run the following command to switch the directory:

```
cd ${BIGDATA_HOME}/om-server/om/sbin
```

**Step 3** Run the following command to configure scheduled backup of FusionInsight Manager's alarm and audit information or service audit information:

```
./setNorthBound.sh -t Information type -i Remote server IP address -p SFTP or FTP port used by the server -u Username -d Save path -c Interval (minutes) -m Number of records in each file -s Whether to enable backup -e Protocol
```

Example:

- `./setNorthBound.sh -t alarm -i 10.0.0.10 -p 22 -u sftpuser -d /tmp/ -c 10 -m 100 -s true -e sftp`

This script modifies the alarm backup configuration file **alarm\_collect\_upload.properties**. The file save path is **{BIGDATA\_HOME}/om-server/tomcat/webapps/web/WEB-INF/classes/config**.

- **./setNorthBound.sh -t audit -i 10.0.0.10 -p 22 -u sftpuser -d /tmp/ -c 10 -m 100 -s true -e sftp**

This script modifies the audit backup configuration file **audit\_collect\_upload.properties**. The file save path is **{BIGDATA\_HOME}/om-server/tomcat/webapps/web/WEB-INF/classes/config**.

- **./setNorthBound.sh -t service\_audit -i 10.0.0.10 -p 22 -u sftpuser -d /tmp/ -c 10 -m 100 -s true -e sftp**

This script modifies the service audit backup configuration file **service\_audit\_collect\_upload.properties**. The file save path is **{BIGDATA\_HOME}/om-server/tomcat/webapps/web/WEB-INF/classes/config**.

- Step 4** Enter the password as prompted. The password is encrypted and saved in the configuration file.

```
Please input sftp/ftp server password:
```

- Step 5** Check the configuration result. If the following information is displayed, the configuration is successful. The configuration file will be automatically synchronized to the standby management node.

```
execute command syncfile successfully.
Config Succeed.
```

----End

## 7.11.8 Enabling the MRS Cluster Maintenance Mode to Disable Alarm Reporting

FusionInsight Manager allows you to set clusters, services, hosts, or OMSs to the maintenance mode. Objects in maintenance mode do not report alarms. This prevents the system from generating a large number of unnecessary alarms during maintenance changes, such as upgrade, because these alarms may influence O&M personnel's judgment on the cluster status.

- **Cluster maintenance mode**  
If a cluster is not brought online or has been brought offline due to O&M operations (for example, non-rolling upgrade), you can set the entire cluster to the maintenance mode.
- **Service maintenance mode**  
When performing maintenance operations on a specific service (for example, performing service-affecting commissioning operations like batch restart of service instances, directly powering on or off nodes of the service, or repairing the service), you can set only this service to the maintenance mode.
- **Host maintenance mode**  
When performing maintenance operations on a host (such as powering on or off, isolating, or reinstalling the host, upgrading its OS, or replacing the host), you can set only this host to the maintenance mode.



- OMS maintenance mode  
When restarting, replacing, or repairing an OMS node, you can set the OMS node to the maintenance mode.

 **NOTE**

This section applies only to MRS 3.x or later.

## Impact on the System

After the maintenance mode is set, alarms caused by non-maintenance operations are suppressed and cannot be reported. Alarms can be reported only when faults persist after the system exits the maintenance mode. Therefore, exercise caution when setting the maintenance mode.



## Enabling the MRS Cluster Maintenance Mode



**Step 1** Log in to FusionInsight Manager.

**Step 2** Set the maintenance mode.

Determine the object to set the maintenance mode based on the service scenario. For details, see [Table 7-98](#).

**Table 7-98** Setting to the maintenance mode

| Scenario                                           | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure a cluster to enter the maintenance mode. | <ol style="list-style-type: none"> <li>1. On the management page, choose <b>***</b> or <b>More &gt; Enter Maintenance</b>.</li> <li>2. In the displayed dialog box, click <b>OK</b>.<br/>After the cluster enters the maintenance state, the status of the cluster becomes . After maintenance is complete, click <b>Exit Maintenance Mode</b>. The cluster then exits the maintenance mode.</li> </ol>                                                                                                                                                                                                                                                                                                                     |
| Configure a service to enter the maintenance mode. | <ol style="list-style-type: none"> <li>1. On FusionInsight Manager, choose <b>Cluster &gt; Services &gt; Service name</b>.</li> <li>2. On the service details page, click <b>More</b> and select <b>Enter Maintenance Mode</b>.</li> <li>3. In the displayed dialog box, click <b>OK</b>.<br/>After a service enters the maintenance mode, the status of the service becomes  in the service list. After maintenance is complete, click <b>Exit Maintenance Mode</b>. The service then exits the maintenance mode.</li> </ol> <p><b>NOTE</b><br/>When configuring a service to enter the maintenance mode, you are advised to set the upper-layer services that depend on this service to the maintenance mode as well.</p> |

| Scenario                                         | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure a host to the maintenance mode.        | <ol style="list-style-type: none"> <li>1. On FusionInsight Manager, choose <b>Hosts</b>.</li> <li>2. On the <b>Hosts</b> page, select the target host, click <b>More</b>, and select <b>Enter Maintenance Mode</b>.</li> <li>3. In the displayed dialog box, click <b>OK</b>.<br/>After the host enters the maintenance mode, the status of the host becomes  in the host list. After maintenance is complete, click <b>Exit Maintenance Mode</b>. The host then exits the maintenance mode.</li> </ol> |
| Configure the OMS to enter the maintenance mode. | <ol style="list-style-type: none"> <li>1. On FusionInsight Manager, choose <b>System &gt; OMS &gt; Enter Maintenance Mode</b>.</li> <li>2. In the displayed dialog box, click <b>OK</b>.<br/>After the OMS enters the maintenance state, the OMS status becomes . After maintenance is complete, click <b>Exit Maintenance Mode</b>. The OMS then exits the maintenance mode.</li> </ol>                                                                                                                |


**Step 3** Check the cluster maintenance view.

On FusionInsight Manager, click **\*\*\*** or **More** next to the cluster name and select **Maintenance Mode View**. In the displayed window, you can view the services and hosts in maintenance mode in the cluster.

After maintenance is complete, you can select services and hosts in batches in the maintenance mode view and click **Exit Maintenance Mode** to make them exit the maintenance mode.

**Figure 7-50** Exiting the maintenance mode in batches

## Maintenance Mode View



----End

## 7.11.9 Configuring Notifications for MRS Cluster Alarms and Events

MRS uses SMN to offer a publish/subscribe model to achieve one-to-multiple message subscriptions and notifications in a variety of message types (SMSs and emails).

On the MRS management console, you can enable or disable the notification service on the **Alarms** page. The functions in the following scenarios can be implemented only after the required cluster function is enabled:

- After a user subscribes to the notification service, the MRS management plane notifies the user of success or failure of manual cluster scale-out and scale-in, cluster deletion, and auto scaling by emails or SMS messages.
- The management plane checks the alarms about the MRS cluster and sends a notification to the tenant if the alarms are critical.
- If either of the operations such as deletion, shutdown, specifications modification, restart, and OS update is performed on an ECS in a cluster, the MRS cluster works abnormally. The management plane notifies a user when detecting that the VM of the user is in either of the preceding operations.

### Creating a Topic

A topic is a specified event for message publication and notification subscription. It serves as a message sending channel, where publishers and subscribers can interact with each other.

- Step 1** Log in to the management console.
  - Step 2** Click **Service List**. Under **Management & Governance**, click **Simple Message Notification**.  
The **SMN** page is displayed.
  - Step 3** In the navigation pane, choose **Topic Management > Topics**.  
The **Topics** page is displayed.
  - Step 4** Click **Create Topic**.  
The **Create Topic** dialog box is displayed.
  - Step 5** In **Topic Name**, enter a topic name. In **Display Name**, enter a display name.
  - Step 6** Select an existing project from the **Enterprise Project** drop-down list, or click **Create Enterprise Project** to create an enterprise project on the **Enterprise Project Management** page and then select it.
  - Step 7** Set tag keys and tag values. Tags consist of keys and values. They identify cloud resources so that you can easily categorize and search for your resources.
- End

## Adding Subscriptions to a Topic

To deliver messages published to a topic to subscribers, you must add subscription endpoints to the topic. SMN automatically sends a confirmation message to the subscription endpoint. The confirmation message is valid only within 48 hours. The subscribers must confirm the subscription within 48 hours so that they can receive notification messages. Otherwise, the confirmation message becomes invalid, and you need to send it again.

- Step 1** Log in to the management console.
  - Step 2** Under **Management & Governance**, click **Simple Message Notification**.  
The **SMN** page is displayed.
  - Step 3** In the navigation pane, choose **Topic Management > Topics**.  
The **Topics** page is displayed.
  - Step 4** Locate the topic to which you want to add a subscription, click **More** in the **Operation** column, and select **Add Subscription**.  
The **Add Subscription** box is displayed.  
**Endpoint** indicates the address of the subscription endpoint. SMS and email, endpoints can be entered in batches. When adding endpoints in batches, each endpoint address occupies a line. You can enter a maximum of 10 endpoints.
  - Step 5** Click **OK**.  
The subscription you added is displayed in the subscription list.
- End

## Sending Notifications to Subscribers

- Step 1** Log in to the MRS console.
- Step 2** Choose **Active Clusters**, select a running cluster, and click its name to go to its details page.
- Step 3** Click **Alarms**.
- Step 4** Choose **Notification Rules > Add Notification Rule**.

**Figure 7-51** Creating a message subscription rule

- Step 5** Set the notification rule parameters.

**Table 7-99** Parameters of a notification rule

| Parameter            | Description                                                                                                                                                                                                                                                                                |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rule Name            | User-defined notification rule name. Only digits, letters, hyphens (-), and underscores (_) are allowed.                                                                                                                                                                                   |
| Message Notification | <ul style="list-style-type: none"> <li>• If you enable this function, the system sends notifications to subscribers based on the notification rule.</li> <li>• If you disable this function, the rule does not take effect, that is, notifications are not sent to subscribers.</li> </ul> |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Topic Name         | Select an existing topic or click <b>Create Topic</b> to create a topic.                                                                                                                                                                                                                                                                                      |
| Notification Type  | Select the type of the notification to be subscribed to. <ul style="list-style-type: none"><li>• Alarm</li><li>• Event</li></ul>                                                                                                                                                                                                                              |
| Subscription Items | Select the items to be subscribed to. You can select all or some items as required.<br>Subscription rules in MRS 3.x or later:<br>Alarm severity: critical, major, minor, and suggestion<br>Event: major, minor, and warning<br>Subscription rules in versions earlier than MRS 3.x:<br>Alarm: critical, major, and minor<br>Event: major, minor, and warning |

**Step 6** Click **OK**.

 **NOTE**

After a message subscription rule is applied, you may receive some historical alarms.

----End

## 7.12 MRS Cluster Alarm Handling Reference

### 7.12.1 ALM-12001 Audit Log Dumping Failure

#### Description

Cluster audit logs need to be dumped on a third-party server due to the local historical data backup policy. The system starts to check the dump server at 3 a.m. every day. If the dump server meets the configuration conditions, audit logs can be successfully dumped. This alarm is generated when the audit log dump fails if the disk space of the dump directory on the third-party server is insufficient or a user changes the username, password, or dump directory of the dump server.

#### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12001    | Minor          | Yes        |

## Parameters

| Name        | Meaning                                                           |
|-------------|-------------------------------------------------------------------|
| Source      | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated.           |
| RoleName    | Specifies the role for which the alarm is generated.              |
| HostName    | Specifies the host for which the alarm is generated.              |

## Impact on the System

System can store a maximum of only 50 dump files locally. If the fault persists on the dump server, the local audit logs may be lost, and the first 50 audit logs that exceed the current time cannot be queried.

## Possible Causes

- The network connection is abnormal.
- The username, password, or dump directory of the dump server does not meet the configuration conditions.
- The disk space of the dump directory is insufficient.

## Procedure

**Check whether the network connection is normal.**

**Step 1** On the FusionInsight Manager home page, choose **Audit > Configurations**.

**Step 2** Check whether the SFTP IP on the dump configuration page is valid.

Log in to the node where Manager is located as user **root** and run the **ping** command to check whether the network connection between the SFTP server and the cluster is normal.

- If yes, go to [Step 5](#).
- If no, go to [Step 3](#).

**Step 3** Repair the network connection, reset the SFTP password, and click **OK**.

**Step 4** Wait for 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Check whether the username, password, or dump directory are correct.**

**Step 5** On the dump configuration page, check whether the username, password, and dump directory of the third-party server are correct.

- If yes, go to [Step 8](#).
- If no, go to [Step 6](#).

**Step 6** Change the username, password, or dump directory, reset the SFTP password and click **OK**.

**Step 7** Wait for 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Check whether the disk space of the dump directory is sufficient.**

**Step 8** Log in to the third-party server as user **root** and run the **df** command to check whether the disk space of the dump directory of the third-party server exceeds 100 MB.

- If yes, go to [Step 11](#).
- If no, go to [Step 9](#).

**Step 9** Expand disk space capacity for the third-party server, Reset the SFTP password and click **OK**

**Step 10** Wait for 2 minutes, view real-time alarms and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Reset the dump rule.**

**Step 11** On the FusionInsight Manager home page, choose **Audit > Configurations**.

**Step 12** Reset dump rules, set the parameters properly, and click **OK**.


**Step 13** Wait for 2 minutes, view real-time alarms and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 14](#).

**Collect fault information.**

**Step 14** On the FusionInsight Manager, choose **O&M > Log > Download**.

**Step 15** Select **OmmServer** from the **Service** and click **OK**.

**Step 16** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 17** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None



## 7.12.2 ALM-12004 OLdap Resource Abnormal

### Description

The system checks LDAP resources every 60 seconds. This alarm is generated when the system detects that the LDAP resources in Manager are abnormal for six consecutive times.

This alarm is cleared when the Ldap resource in the Manager recovers and the alarm handling is complete.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12004    | Major          | Yes        |

### Parameters

| Name        | Meaning                                                           |
|-------------|-------------------------------------------------------------------|
| Source      | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated.           |
| RoleName    | Specifies the role for which the alarm is generated.              |
| HostName    | Specifies the host for which the alarm is generated.              |

### Impact on the System

Manager active/standby switchover may occur. The Manager and WebUI authentication service of components is unavailable. Security authentication and user management functions cannot be provided for upper-layer web services. As a result, you may fail to log in to the Manager and WebUI of components.

### Possible Causes

The LdapServer process in the Manager is abnormal.

### Procedure

**Check whether the LdapServer process in the Manager is normal.**

**Step 1** Log in the Manager node in the cluster as user **omm**.

Log in to FusionInsight Manager using the floating IP address, and run the **sh \${BIGDATA\_HOME}/om-server/om/sbin/status-oms.sh** command to check the information about the current Manager two-node cluster.

**Step 2** Run **ps -ef | grep slapd** command to check whether the LdapServer resource process in the **\${BIGDATA\_HOME}/om-server/om/** in the process configuration file is running properly.

 **NOTE**

You can determine that the resource is normal by checking the following information:

1. After the **sh \${BIGDATA\_HOME}/om-server/om/sbin/status-oms.sh** command runs, **ResHAStatus** of the OLdap is **Normal**.
2. After the **ps -ef | grep slapd** command runs, the slapd process of port 21750 can be viewed.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 4](#).


**Step 3** Run the **kill -2 ldap pid** command to restart the LdapServer process and wait for 20 seconds. The HA starts the OLdap process automatically. Check whether the current OLdap resource is in normal state.

- If yes, the operation is complete.
- If no, go to [Step 4](#).

**Collect fault information.**

**Step 4** On the FusionInsight Manager home page, choose **O&M > Log > Download**.

**Step 5** Select **OmsLdapServer** and **OmmServer** from the **Service** and click **OK**.

**Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.3 ALM-12005 OKerberos Resource Abnormal

### Description

The alarm module checks the status of the Kerberos resource in Manager every 80 seconds. This alarm is generated when the alarm module detects that the Kerberos resources are abnormal for six consecutive times.

This alarm is cleared when the Kerberos resource recovers and the alarm handling is complete.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12005    | Major          | Yes        |

## Parameters

| Name        | Meaning                                                           |
|-------------|-------------------------------------------------------------------|
| Source      | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated.           |
| RoleName    | Specifies the role for which the alarm is generated.              |
| HostName    | Specifies the host for which the alarm is generated.              |

## Impact on the System

The component WebUI authentication services are unavailable and cannot provide security authentication functions for web upper-layer services. Users may be unable to log in to FusionInsight Manager and the WebUIs of components.

## Possible Causes

The OLdap resource on which the Okerberos depends is abnormal.

## Procedure

**Check whether the OLdap resource on which the Okerberos depends is abnormal in the Manager.**

**Step 1** Log in the Manager node in the cluster as user **omm**.

Log in to FusionInsight Manager using the floating IP address, and run the **sh \$ {BIGDATA\_HOME}/om-server/om/sbin/status-oms.sh** command to check the information about the current Manager two-node cluster.

**Step 2** Run the **sh \$ {BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** command to check whether the OLdap resource status managed by HA is normal. (In single-node mode, the OLdap resource is in the Active\_normal state; in the two-node mode, the OLdap resource is in the Active\_normal state on the active node and in the Standby\_normal state on the standby node.)

- If yes, go to [Step 4](#).

- If no, go to [Step 3](#).


**Step 3** See the procedure in [ALM-12004 OLdap Resource Abnormal](#) to resolve the problem. After the OLdap resource status recovers, check whether the OKerberos resource status is normal.

- If yes, the operation is complete.
- If no, go to [Step 4](#).

**Collect fault information.**

**Step 4** On the FusionInsight Manager home page, choose **O&M > Log > Download**.

**Step 5** Select **OmsKerberos** and **OmmServer** from the **Service** and click **OK**.

**Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.4 ALM-12006 Node Fault

### Alarm Description

Controller checks the NodeAgent heartbeat every 30 seconds. If Controller does not receive heartbeat messages from a NodeAgent, it attempts to restart the NodeAgent process. This alarm is generated if the NodeAgent fails to be restarted for three consecutive times.

This alarm is cleared when Controller can properly receive the status report of the NodeAgent.

#### NOTE

In MRS 3.3.0 and later versions, the alarm name is changed to **NodeAgent Process Is Abnormal**.

### Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 12006    | Major          | Yes          |

## Alarm Parameters

| Parameter   | Description                                                        |
|-------------|--------------------------------------------------------------------|
| Source      | Specifies the cluster or system for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated.           |
| RoleName    | Specifies the role for which the alarm was generated.              |
| HostName    | Specifies the host for which the alarm was generated.              |

## Impact on the System


NodeAgent process is abnormal, heartbeat messages cannot be reported to the platform. If the problem is caused by network faults, hardware faults, or SSH mutual trust, component services cannot be normal.

## Possible Causes

- The network is disconnected, the hardware is faulty, or the operating system runs slowly.
- The memory of the NodeAgent process is insufficient.
- The NodeAgent process is faulty.

## Handling Procedure

**Check whether the network is disconnected, whether the hardware is faulty, or whether the operating system runs commands slowly.**

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the page that is displayed, click  in the row containing the alarm, click the host name, and view the IP address of the host for which the alarm is generated.

**Step 2** Log in to the active management node as user **root**.

### NOTE

If the faulty node is the active management node and fails login, the network of the active management node may be faulty. In this case, go to [Step 4](#).

**Step 3** Run the **ping *IP address of the faulty host*** command to check whether the faulty node is reachable.

- If yes, go to [Step 12](#).
- If no, go to [Step 4](#).

**Step 4** Contact the network administrator to check whether the network is faulty.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Step 6** Contact the hardware administrator to check whether the hardware (CPU or memory) of the node is faulty.

- If yes, go to [Step 7](#).
- If no, go to [Step 12](#).

**Step 7** Repair or replace faulty components and restart the node. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Step 8** If a large number of node faults are reported in the cluster, the floating IP addresses may be abnormal. As a result, Controller cannot detect the NodeAgent heartbeat.

Log in to any management node and view the `/var/log/Bigdata/omm/oms/ha/scriptlog/floatip.log` log to check whether the logs generated one to two minutes before and after the faults occur are complete.

For example, a complete log is in the following format:

```
2017-12-09 04:10:51,000 INFO (floatip) Read from ${BIGDATA_HOME}/om-server_*/om/etc/om/routeSetConf.ini,value is : yes
2017-12-09 04:10:51,000 INFO (floatip) check wsNetExport : eth0 is up.
2017-12-09 04:10:51,000 INFO (floatip) check omNetExport : eth0 is up.
2017-12-09 04:10:51,000 INFO (floatip) check wsInterface : eRth0:oms, wsFloatIp: XXX.XXX.XXX.XXX.
2017-12-09 04:10:51,000 INFO (floatip) check omInterface : eth0:oms, omFloatIp: XXX.XXX.XXX.XXX.
2017-12-09 04:10:51,000 INFO (floatip) check wsFloatIp : XXX.XXX.XXX.XXX is reachable.
2017-12-09 04:10:52,000 INFO (floatip) check omFloatIp : XXX.XXX.XXX.XXX is reachable.
```

- If yes, go to [Step 12](#).
- If no, go to [Step 9](#).

**Step 9** Check whether the omNetExport log is printed after the wsNetExport is detected or whether the interval for printing two logs exceeds 10 seconds or longer.

- If yes, go to [Step 10](#).
- If no, go to [Step 12](#).

**Step 10** View the `/var/log/message` file of the OS to check whether sssd frequently restarts or nscd exception information is displayed when the fault occurs.

sssd restart example

```
Feb 7 11:38:16 10-132-190-105 sssd[pam]: Shutting down
Feb 7 11:38:16 10-132-190-105 sssd[nss]: Shutting down
Feb 7 11:38:16 10-132-190-105 sssd[nss]: Shutting down
Feb 7 11:38:16 10-132-190-105 sssd[be[default]]: Shutting down
Feb 7 11:38:16 10-132-190-105 sssd: Starting up
Feb 7 11:38:16 10-132-190-105 sssd[be[default]]: Starting up
Feb 7 11:38:16 10-132-190-105 sssd[nss]: Starting up
Feb 7 11:38:16 10-132-190-105 sssd[pam]: Starting up
```

Example nscd exception information

```
Feb 11 11:44:42 10-120-205-33 nscd: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.55:21780:
Can't contact LDAP server
```

```
Feb 11 11:44:43 10-120-205-33 ntpq: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.55:21780:
Can't contact LDAP server
Feb 11 11:44:44 10-120-205-33 ntpq: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.92:21780:
Can't contact LDAP server
```

- If yes, go to [Step 11](#).
- If no, go to [Step 12](#).

**Step 11** Check whether the LdapServer node is faulty, for example, the service IP address is unreachable or the network latency is too high. If the fault occurs periodically, locate and eliminate it and run the **top** command to check whether abnormal software exists.

**Check whether the memory of the NodeAgent process is insufficient.**

**Step 12** Log in to the faulty node as user **root** and run the following command to view the NodeAgent process logs:

```
vi /var/log/Bigdata/nodeagent/scriptlog/agent_gc.log.*.current
```

**Step 13** Check whether the log file contains an error indicating that the metaspace size or heap memory size is insufficient.

- If yes, go to [Step 14](#).
- If no, go to [Step 17](#).

**Step 14** Run the **su - omm** command to switch to user **omm**, edit the corresponding file based on the cluster version, increase the values of **nodeagent.Xms** (initial heap memory) and **nodeagent.Xmx** (maximum heap memory), and save the modification.

The path of the file containing the parameters is as follows:

- Versions earlier than MRS 3.2.1: **/opt/Bigdata/om-agent/nodeagent/bin/nodeagent\_ctl.sh**
- MRS 3.2.1 or later: **\$NODE\_AGENT\_HOME/etc/agent/nodeagent.properties**

**Step 15** Run the following commands to restart the NodeAgent service:

```
sh ${BIGDATA_HOME}/om-agent/nodeagent/bin/stop-agent.sh
```

```
sh ${BIGDATA_HOME}/om-agent/nodeagent/bin/start-agent.sh
```

**Step 16** Wait a moment and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 17](#).

**Check whether the NodeAgent process is faulty.**

**Step 17** Log in to the faulty node as user **omm** and run the following command:

```
ps -ef | grep "Dprocess.name=nodeagent" | grep -v grep
```

**Step 18** Check whether the command output is empty.

- If yes, go to [Step 19](#).
- If no, go to [Step 21](#).

**Step 19** View the NodeAgent startup and run logs to locate the fault. After the fault is rectified, go to [Step 20](#).

- NodeAgent run logs: `/var/log/Bigdata/nodeagent/agentlog/agent.log`
- NodeAgent start and stop logs: `/var/log/Bigdata/nodeagent/scriptlog/nodeagent_ctl.log`

**Step 20** Run the following commands to restart the NodeAgent service:

```
sh ${BIGDATA_HOME}/om-agent/nodeagent/bin/stop-agent.sh
```


```
sh ${BIGDATA_HOME}/om-agent/nodeagent/bin/start-agent.sh
```

**Collect fault information.**

**Step 21** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 22** Select the following nodes from **Services** and click **OK**.

- NodeAgent
- Controller
- OS

**Step 23** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 24** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.5 ALM-12007 Process Fault

### Description

This alarm is generated when the process health check module detects that the process connection status is **Bad** for three consecutive times. The process health check module checks the process status every 5 seconds.

This alarm is cleared when the process can be connected.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12007    | Major          | Yes        |



## Parameters

| Name        | Meaning                                                           |
|-------------|-------------------------------------------------------------------|
| Source      | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated.           |
| RoleName    | Specifies the role for which the alarm is generated.              |
| HostName    | Specifies the host for which the alarm is generated.              |

## Impact on the System

The impact varies depending on the instance that is faulty.

For example, if an HDFS instance is faulty, the impacts are as follows:

- If a DataNode instance is faulty, read and write operations cannot be performed on data blocks stored on the DataNode, which may cause data loss or unavailability. However, data in HDFS is redundant. Therefore, the client can access data from other DataNodes.
- If an HttpFS instance is faulty, the client cannot access files in HDFS over HTTP. However, the client can use other methods (such as shell commands) to access files in HDFS.
- If a JournalNode instance is faulty, namespaces and data logs cannot be stored to disks, which may cause data loss or unavailability. However, HDFS stores backups on other JournalNodes. Therefore, the faulty JournalNode can be recovered and data can be rebalanced.
- If a NameNode deployed in active/standby mode is faulty, an active/standby switchover occurs. If only one NameNode is deployed, the client cannot read or write any HDFS data. On MRS, NameNodes must be deployed in two-node mode.
- If a Router instance is faulty, the client cannot access data on the router. However, the client can use other Routers or directly access data on the backend NameNode.
- If a ZKFC instance is faulty, the NameNode does not continuously and automatically fail over. As a result, data cannot be read from or write to HDFS by the client. In this case, you need to enable automatic failover on other available ZKFC instances to restore the HDFS cluster.

## Possible Causes


- The instance process is abnormal.
- The disk space is insufficient.

 NOTE

If a large number of process fault alarms exist in a time segment, files in the installation directory may be deleted mistakenly or permission on the directory may be modified.

## Procedure

### Check whether the instance process is abnormal.

**Step 1** In the FusionInsight Manager portal, click **O&M > Alarm > Alarms**, click  in the row where the alarm is located, and click the host name to view the host address for which the alarm is generated

**Step 2** On the **Alarms** page, check whether the **ALM-12006 Node Fault** is generated.

- If yes, go to [Step 3](#).
- If no, go to [Step 4](#).

**Step 3** Handle the alarm according to **ALM-12006 Node Fault**.

**Step 4** Log in to the host for which the alarm is generated as user **root**. Check whether the installation directory user, user group, and permission of the alarm role are correct. The user, user group, and the permission must be **omm:ficommon 750**.

For example, the NameNode installation directory is `${BIGDATA_HOME}/FusionInsight_Current/1_8_NameNode/etc`.

- If yes, go to [Step 6](#).
- If no, go to [Step 5](#).

**Step 5** Run the following command to set the permission to **750** and **User:Group** to **omm:ficommon**:

```
chmod 750 <folder_name>
```

```
chown omm:ficommon <folder_name>
```

**Step 6** Wait for 5 minutes. In the alarm list, check whether **ALM-12007 Process Fault** is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Step 7** Log in to the active OMS node as user **root** and run the following command to view the **configurations.xml** file. In the preceding command, "Service name" is the service name queried in [Step 1](#).

```
vi $BIGDATA_HOME/components/current/Service name/configurations.xml
```

Search for the keyword **healthMonitor.properties**, find the health check configuration item corresponding to the alarm reporting instance, and record the interface or script path specified by **monitor.info**, as shown in the following figure.

Check the logs recorded in the interface or script and rectify the fault.

```
<config category="healthMonitor.properties" format="propertyfileconfigurer">
 <property type="hidden" scope="setup" classification="System">
 <name>monitor.type</name>
 <value>SCRIPT</value>
 </property>
 <property type="hidden" scope="setup" classification="System">
 <name>monitor.preInitDelay</name>
 <value>120000</value>
 <!-- 2min -->
 </property>
 <property type="hidden" scope="setup" classification="System">
 <name>monitor.recheckTimes</name>
 <value>90</value>
 </property>
 <property type="hidden" scope="setup" classification="System">
 <name>monitor.checkIntervals</name>
 <value>10000</value>
 </property>
 <property type="hidden" scope="setup" classification="System">
 <name>monitor.info</name>
 <value>#{install_home}/sbin/dbservice_mon.sh</value>
 </property>
 <property type="hidden" scope="setup" classification="System">
 <name>monitor.metric.connector</name>
 <value>#{install_home}/sbin/dbservice_ha.sh</value>
 </property>
</config>
```

**Step 8** Wait for 5 minutes. In the alarm list, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Check whether disk space is sufficient.**

**Step 9** On the FusionInsight Manager, check whether the alarm list contains **ALM-12017 Insufficient Disk Capacity**.

- If yes, go to [Step 10](#).
- If no, go to [Step 13](#).

**Step 10** Rectify the fault by following the steps provided in **ALM-12017 Insufficient Disk Capacity**.

**Step 11** Wait for 5 minutes. In the alarm list, check whether **ALM-12017 Insufficient Disk Capacity** is cleared.

- If yes, go to [Step 12](#).
- If no, go to [Step 13](#).


**Step 12** Wait for 5 minutes. In the alarm list, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 13](#).

**Collect fault information.**

**Step 13** On the FusionInsight Manager, choose **O&M > Log > Download**.

**Step 14** According to the service name obtained in [Step 1](#), select the component and **NodeAgent** from the **Service** and click **OK**.

**Step 15** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 16** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.6 ALM-12010 Manager Heartbeat Interruption Between the Active and Standby Nodes

### Description

This alarm is generated when the active Mager does not receive the heartbeat signal from the standby Manager within 7 seconds.

This alarm is cleared when the active Manager receives heartbeat signals from the standby Manager.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12010	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

When the active Manager process is abnormal, the active/standby switchover cannot be performed, affecting basic O&M functions.


### Possible Causes

- The link between the active and standby Manager is abnormal.

- The node name configuration is incorrect.
- The port is disabled by the firewall.

## Procedure

**Check whether the network between the active and standby Manager server is normal.**

- Step 1** In the FusionInsight Manager portal, click **O&M > Alarm > Alarms**, click  in the row containing the alarm and view the IP address of the standby Manager (Peer Manager) server in the alarm details.
- Step 2** Log in to the active OMS node as user **root**.
- Step 3** Run the **ping standby Manager heartbeat IP address** command to check whether the standby Manager server is reachable.
- If yes, go to [Step 6](#).
  - If no, go to [Step 4](#).
- Step 4** Contact the network administrator to check whether the network is faulty.
- If yes, go to [Step 5](#).
  - If no, go to [Step 6](#).
- Step 5** Rectify the network fault and check whether the alarm is cleared from the alarm list.
- If yes, no further action is required.
  - If no, go to [Step 6](#).

**Check whether the node name is correctly configured.**

- Step 6** Run the following command to go to the software installation directory of the active OMS node:
- ```
cd /opt
```
- Step 7** Run the following command to find the configuration file directory of the active and standby nodes.
- ```
find -name hacom_local.xml
```
- Step 8** Run the following command to go to the **workspace** directory:
- ```
cd${BIGDATA_HOME}/om-server/OMS/workspace0/ha/local/hacom/conf/
```
- Step 9** Run the **vim** command to open the **hacom_local.xml** file. Check whether the local and peer nodes are correctly configured. The local node is configured as the active node, and the peer node is configured as the standby node.
- If yes, go to [Step 12](#).
 - If no, go to [Step 10](#).
- Step 10** Modify the configuration of the active and standby nodes in the **hacom_local.xml** file and press **Esc** to return to the command mode. Run the **:wq** command to save the modification and exit.
- Step 11** Check whether the alarm is cleared automatically.

- If yes, no further action is required.
- If no, go to [Step 12](#).

Check whether the port is disabled by the firewall.

Step 12 Run the `lsof -i :20012` command to check whether the heartbeat ports of the active and standby nodes are enabled. If the command output is displayed, the ports are enabled. Otherwise, the ports are disabled by the firewall.

- If yes, go to [Step 13](#).
- If no, go to [Step 16](#).

Step 13 Run the `iptables -P INPUT ACCEPT` command to avoid the server disconnection.

Step 14 Run the following command to clear the firewall:

`iptables -F`

Step 15 Check whether the alarm is cleared from the alarm list.


- If yes, no further action is required.
- If no, go to [Step 16](#).

Collect fault information.

Step 16 On the FusionInsight Manager, choose **O&M > Log > Download**.

Step 17 Select the following nodes from the **Service** and click **OK**:

- OmmServer
- Controller
- NodeAgent

Step 18 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.7 ALM-12011 Manager Data Synchronization Exception Between the Active and Standby Nodes

Description

The system checks data synchronization between the active and standby Manager nodes every 60 seconds. This alarm is generated when the standby Manager fails to synchronize files with the active Manager.

This alarm is cleared when the standby Manager synchronizes files with the active Manager.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12011 | Critical | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System


The configuration file of the standby Manager is not updated. When an active/standby switchover occurs, the configuration file that fails to be synchronized may be lost. Manager and some components may not run properly.

Possible Causes

- The link between the active and standby Managers is interrupted or The storage space of the **/srv/BigData/LocalBackup** directory is full.
- The synchronization file does not exist or the file permission is incorrect.

Procedure

Check whether the network between the active Manager server and the standby Manager server is normal.

- Step 1** In the FusionInsight Manager portal, click **O&M > Alarm > Alarms**, click  in the row where the alarm is located and obtain the standby Manager server IP address (Peer Manager IP address) in the alarm details.
- Step 2** Log in to the active Manager server as user **root**.
- Step 3** Run the **ping *standby Manager IP address*** command to check whether the standby Manager server is reachable.
 - If yes, go to **Step 6**.

- If no, go to [Step 4](#).

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Check whether the storage space of the `/srv/BigData/LocalBackup` directory is full.

Step 6 Run the following command to check whether the storage space of the `/srv/BigData/LocalBackup` directory is full:

```
df -hl /srv/BigData/LocalBackup
```

- If yes, go to [Step 7](#).
- If no, go to [Step 10](#).

Step 7 Run the following command to clear unnecessary backup files:

```
rm -rf Directory to be cleared
```

Example:

```
rm -rf /srv/BigData/LocalBackup/0/default-oms_20191211143443
```

Step 8 On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

In the **Operation** column of the backup task to be performed, click **Configure** and change the value of **Maximum Number of Backup Copies** to reduce the number of backup file sets.

Step 9 Wait about 1 minute and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

Check whether the synchronization file exists and whether the file permission is normal.

Step 10 Run the following command to check whether the synchronization file exists.

```
find /srv/BigData/ -name "sed*"
```

```
find /opt -name "sed*"
```

- If yes, go to [Step 11](#).
- If no, go to [Step 12](#).

Step 11 Run the following command to view the synchronization file information and permission obtained in [Step 10](#).

```
ll path of the file to be found
```


- If the size of the file is 0 and the permission column is -, the file is a junk file. Run the following command to delete it.

```
rm -rf files to be deleted
```

Wait for several minutes and check whether the alarm is cleared. If the alarm persists, go to [Step 12](#).

- If the file size is not 0, go to [Step 12](#).

Step 12 View the log files generated when the alarm is generated.

1. Run the following command to switch to the HA run log file path.

```
cd /var/log/Bigdata/omm/oms/ha/runlog/
```

2. Decompress and view the log files generated when the alarm is generated.

For example, if the name of the file to be viewed is

ha.log.2021-03-22_12-00-07.gz, run the following command:

```
gunzip ha.log.2021-03-22_12-00-07.gz
```

```
vi ha.log.2021-03-22_12-00-07
```

Check whether error information is reported before and after the alarm generation time.

- If yes, rectify the fault based on the error information. Then go to [Step 13](#).

For example, if the following error information is displayed, the directory permission is insufficient. In this case, change the directory permission to be the same as that on the normal node.

- If no, go to [Step 14](#).

Step 13 Wait about 10 minute and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 14](#).

Collect fault information.

Step 14 On the FusionInsight Manager, choose **O&M > Log > Download**.

Step 15 Select the following nodes from the **Service** and click **OK**:

- OmmServer
- Controller
- NodeAgent

Step 16 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 17 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.8 ALM-12012 NTP Service Is Abnormal

Alarm Description

The system checks whether the NTP service on a node synchronizes time with the NTP service on the active OMS node every 60 seconds. This alarm is generated when the NTP service fails to synchronize time for two consecutive times.

This alarm is generated when the time difference between the NTP service on a node and the NTP service on the active OMS node is greater than or equal to 20s for two consecutive times. This alarm is cleared when the time difference is less than 20s.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 12012 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster or system for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |

Impact on the System

The time on the node is inconsistent with that on other nodes in the cluster. Therefore, some FusionInsight applications on the node may not run properly. If the time difference between the node and other Kerberos service instances keeps increasing, Kerberos authentication on the node may fail and service exceptions occur.

Possible Causes

- The NTP service on the current node cannot start properly.
- The current node fails to synchronize time with the NTP service on the active OMS node.
- The key authenticated by the NTP service on the current node is inconsistent with that on the active OMS node.
- The time offset between the node and the NTP service on the active OMS node is large.

Handling Procedure

Check the NTP service mode of the node.

- Step 1** Log in to the active management node as user **root**, run the **su - omm** command to switch to user **omm**, and run the following command to check the resource status on the active and standby nodes:


```
sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh
```

- If "chrony" is displayed in the **ResName** column of the command output, go to [Step 2](#).
- If "ntp" is displayed in the **ResName** column, go to [Step 20](#).

NOTE

If both "chrony" and "ntp" are displayed in the **ResName** column of the command output, the NTP service mode is being switched. Wait for 10 minutes and go to [Step 1](#) again. If both "chrony" and "ntp" persist, contact O&M personnel.

Check whether the chrony service on the node is started properly.

- Step 2** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the page that is displayed, click  in the row containing the alarm, and view the name of the host for which the alarm is generated in **Location**.
- Step 3** Check whether the chronyd process is running on the node where the alarm is generated. Log in to the node for which the alarm is generated as user **root** and run the **ps -ef | grep chronyd | grep -v grep** command to check whether the command output contains the chronyd process.
- If yes, go to [Step 6](#).
 - If no, go to [Step 4](#).

- Step 4** Start the NTP service.

- Step 5** After 10 minutes, check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 6](#).

Check whether the current node can synchronize time properly with the chrony service on the active OMS node.

- Step 6** Check whether the node can synchronize time with the NTP service on the active OMS node based on additional information of the alarm.
- If yes, go to [Step 7](#).

- If no, go to [Step 17](#).

Step 7 Check whether the synchronization with the chrony service on the active OMS node is faulty.

Log in to the node for which the alarm is generated as user **root** and run the **chronyc sources** command.

In the command output, if there is an asterisk (*) before the IP address of the chrony service on the active OMS node, the synchronization is normal. The command output is as follows:

```
MS Name/IP address      Stratum Poll Reach LastRx Last sample
=====
^* 10.10.10.162         10 10 377 626 +16us[ +15us] +/- 308us
```

In the command output, if there is no asterisk (*) before the IP address of the NTP service on the active OMS node, and the value of **Reach** is **0**, the synchronization is abnormal.

```
MS Name/IP address      Stratum Poll Reach LastRx Last sample
=====
^? 10.1.1.1            0 10 0 - +0ns[ +0ns] +/- 0ns
```

- If yes, go to [Step 8](#).
- If no, go to [Step 38](#).

Step 8 The chrony synchronization failure is typically caused by the system firewall. If the firewall can be disabled, disable it. If the firewall cannot be disabled, check the firewall configuration policy and ensure that UDP ports 123 and 323 are not disabled. (For details, see the firewall configuration policy of each system.)

Step 9 Check whether the alarm is cleared 10 minutes later.

- If yes, no further action is required.
- If no, go to [Step 10](#).

Step 10 Log in to the active OMS node as user **root** and run the following command to view the authentication code whose key index is **1M**:

In Red Hat Enterprise Linux, run the **cat \${BIGDATA_HOME}/om-server/OMS/workspace/conf/chrony.keys** command.

Step 11 Run the following command to check whether the key is the same as that queried in [Step 10](#):

In Red Hat Enterprise Linux, run the **diff \${BIGDATA_HOME}/om-server/OMS/workspace/conf/chrony.keys /etc/chrony.keys** command.

NOTE

If the keys are the same, no result is returned after the command is executed. For example:

```
host01:~ # cat ${BIGDATA_HOME}/om-server/OMS/workspace/conf/chrony.keys
1 M sdYbq;o^CzEAWo<U=Tw5
host01:~ # diff ${BIGDATA_HOME}/om-server/OMS/workspace/conf/chrony.keys /etc/chrony.keys
host01:~ #
```

- If yes, go to [Step 12](#).
- If no, go to [Step 38](#).

Step 12 Run the **cat \${BIGDATA_HOME}/om-server/om/packaged-distributables/ntpKeyFile** command to check whether the key is the same as that queried in

Step 10. (Compare the key with that of the authentication key index **1M** queried in [Step 10](#).)

- If yes, go to [Step 13](#).
- If no, go to [Step 15](#).

Step 13 Log in to the faulty node as user **root** and run the `cat /etc/chrony.keys` command in Red Hat Enterprise Linux to check whether the key is the same as that queried in [Step 12](#) (compare it with that of the authentication key index **1M** for comparison).

- If yes, go to [Step 38](#).
- If no, go to [Step 14](#).

Step 14 Run the `su - omm` command to switch to user **omm**, change the key of the authentication key index **1M** in `/${NODE_AGENT_HOME}/chrony.keys` to the key of `ntpKeyFile` in [Step 12](#), and go to [Step 16](#).

Step 15 Run the following commands as user **root** or **omm** to change the NTP key of the active OMS node (change `ntp.keys` to `ntpkeys` in Red Hat Enterprise Linux):

```
cd ${BIGDATA_HOME}/om-server/OMS/workspace/conf
```

```
sed -i "cat chrony.keys | grep -n '1 M'|awk -F ':' '{print $1}'`d" chrony.keys
```

```
echo "1 M `cat ${BIGDATA_HOME}/om-server/om/packaged-distributables/  
ntpKeyFile`" >> chrony.keys
```

Check whether the key of the authentication key index **1M** in `chrony.keys` is the same as that of `ntpKeyFile`.

- If yes, go to [Step 16](#).
- If no, change the key of the authentication key index **1M** in `chrony.keys` to the key of `ntpKeyFile` and go to [Step 16](#).

Step 16 After 5 minutes, run the `systemctl restart chronyd` command to restart the chrony service on the active OMS node. After 15 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 38](#).

Check whether the time deviation between the node and the chrony service on the active OMS node is large.

Step 17 Check whether the time deviation is large in additional information of the alarm.

- If yes, go to [Step 18](#).
- If no, go to [Step 38](#).

Step 18 On the **Hosts** tab page, select the host for which the alarm is generated, and choose **More > Stop All Instances** to stop all the services on the node.

If the time on the alarm node is later than that on the chrony service of the active OMS node, adjust the time of the alarm node. After adjusting the time, choose **More > Start All Instances** to start the services on the node.

If the time on the alarm node is earlier than that on the chrony service of the active OMS node, wait until the time deviation is due and adjust the time of the

alarm node. After adjusting the time, choose **More > Start All Instances** to start the services on the node.


 **NOTE**

If you do not wait, data loss may occur.

Step 19 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 38](#).

Check whether the NTP service on the node is started properly.

Step 20 On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the page that is displayed, click  in the row containing the alarm, and view the name of the host for which the alarm is generated in **Location**.

Step 21 Check whether the ntpd process is running on the node using the following method. Log in to the alarm node as user **root** and run the **ps -ef | grep ntpd | grep -v grep** command to check whether the command output contains the ntpd process.

- If yes, go to [Step 24](#).
- If no, go to [Step 22](#).

Step 22 Run the **service ntp start** command (or the **service ntpd start** command in Red Hat Enterprise Linux) to start the NTP service.

Step 23 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 24](#).

Check whether the node can synchronize time properly with the NTP service on the active OMS node.

Step 24 Check whether the node can synchronize time with the NTP service on the active OMS node based on additional information of the alarm.

- If yes, go to [Step 25](#).
- If no, go to [Step 35](#).

Step 25 Check whether the synchronization with the NTP service on the active OMS node is faulty.

Log in to the alarm node as user **root** and run the **ntpq -np** command.

If an asterisk (*) exists before the IP address of the NTP service on the active OMS node in the command output, the synchronization is in normal state. The command output is as follows:

```
remote refid st t when poll reach delay offset jitter
=====
*10.10.10.162 .LOCL. 1 u 1 16 377 0.270 -1.562 0.014
```

If there is no asterisk (*) before the IP address of the NTP service on the active OMS node, as shown in the following command output, and the value of **refid** is **.INIT.**, the synchronization is abnormal.

```
remote refid st t when poll reach delay offset jitter
=====
```

```
10.10.10.162 .INIT. 1 u 1 16 377 0.270 -1.562 0.014
```

- If yes, go to [Step 26](#).
- If no, go to [Step 38](#).

Step 26 The NTP synchronization failure is typically caused by the system firewall. If the firewall can be disabled, run the **iptables -F** command to disable it. If the firewall cannot be disabled, run the **iptables -L** command to check the firewall configuration policy and ensure that the UDP port 123 is not disabled. (For details, see the firewall configuration policy of each system.)

Step 27 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 28](#).

Step 28 Log in to the active OMS node as user **root** and run the following command to view the authentication key index **1M**:

In SUSE Linux, run the **cat \${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntp.keys** command.

In Red Hat Enterprise Linux or EulerOS, run the **cat \${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntpkeys** command.

Step 29 Run the following command to check whether the key is the same as that queried in [Step 28](#):

In SUSE Linux, run the **diff \${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntp.keys /etc/ntp.keys** command.

In Red Hat Enterprise Linux or EulerOS, run the **diff \${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntpkeys /etc/ntp/ntpkeys** command.

NOTE

If the keys are the same, no result is returned after the command is executed. For example:

```
host01:~ # cat ${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntp.keys
1 M sdYbq;o^CzEAWo<U=Tw5
host01:~ # diff ${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntp.keys /etc/ntp.keys
host01:~ #
```

- If yes, go to [Step 30](#).
- If no, go to [Step 38](#).

Step 30 Run the **cat \${BIGDATA_HOME}/om-server/om/packaged-distributables/ntpKeyFile** command to check whether the key is the same as that queried in [Step 28](#). (Compare the key with that of the authentication key index **1M** queried in [Step 28](#).)

- If yes, go to [Step 31](#).
- If no, go to [Step 33](#).

Step 31 Log in to the faulty node as user **root** and run the **cat /etc/ntp.keys** command in SUSE Linux (or the **cat /etc/ntp/ntpkeys** command in Red Hat Enterprise Linux) to check whether the key is the same as the value queried in [Step 30](#) (use the key of the authentication key index **1M** for comparison).

- If yes, go to [Step 38](#).
- If no, go to [Step 32](#).

Step 32 Run the `su - omm` command to switch to user `omm`, change the key of the authentication key index `1M` in `${NODE_AGENT_HOME}/ntp.keys` (`${NODE_AGENT_HOME}/ntpkeys` in Red Hat Enterprise Linux) to the key of `ntpKeyFile` in [Step 30](#), and go to [Step 34](#).

Step 33 Run the following commands as user `root` or `omm` to change the NTP key of the active OMS node (change `ntp.keys` to `ntpkeys` in Red Hat Enterprise Linux):

```
cd ${BIGDATA_HOME}/om-server/OMS/workspace/conf
```

```
sed -i "`cat ntp.keys | grep -n '1 M'|awk -F ':' '{print $1}'`d" ntp.keys
```

```
echo "1 M `cat ${BIGDATA_HOME}/om-server/om/packaged-distributables/  
ntpKeyFile`" >>ntp.keys
```

Check whether the key of the authentication key index `1M` in `ntp.keys` is the same as that of `ntpKeyFile`.

- If yes, go to [Step 34](#).
- If no, change the key of the authentication key index `1M` in `ntp.keys` to the key of `ntpKeyFile` and go to [Step 34](#).

Step 34 After 5 minutes, run the `service ntp restart` command to restart the NTP service on the active OMS node. After 15 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 38](#).

Check whether the time deviation between the node and the NTP service on the active OMS node is large.

Step 35 Check whether the time deviation is large in additional information of the alarm.

- If yes, go to [Step 36](#).
- If no, go to [Step 38](#).

Step 36 On the **Hosts** tab page, select the host for which the alarm is generated, and choose **More > Stop All Instances** to stop all the services on the node.

If the time on the alarm node is later than that on the NTP service of the active OMS node, adjust the time of the alarm node. After adjusting the time, choose **More > Start All Instances** to start the services on the node.

If the time on the alarm node is earlier than that on the NTP service of the active OMS node, wait until the time deviation is due and adjust the time of the alarm node. After adjusting the time, choose **More > Start All Instances** to start the services on the node.


NOTE

If you do not wait, data loss may occur.

Step 37 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 38](#).

Collect fault information.

- Step 38** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 39** Expand the **Service** drop-down list, select **NodeAgent** and **OmmServer** for the target cluster, and click **OK**. Expand the **Hosts** dialog box and select the alarm node and the active OMS node.
- Step 40** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 41** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.9 ALM-12014 Partition Lost

Description

The system checks the partition status every 60 seconds. This alarm is generated when the system detects that a partition to which service directories are mounted is lost (because the device is removed or goes offline, or the partition is deleted). The system checks the partition status periodically.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|--|
| 12014 | Major | <ul style="list-style-type: none"> • Yes: MRS 3.3.0 and later, MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions • No: Versions earlier than MRS 3.3.0 |

Parameters

| Name | Meaning |
|---------------|---|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| DirName | Specifies the directory for which the alarm is generated. |
| PartitionName | Specifies the device partition for which the alarm is generated. |


Impact on the System

- Data loss: The device partition is lost and the data stored in the partition is lost.
- System breakdown: If the system disk is lost, the system deployed on the node cannot run properly. In some cases, the system may break down and cannot be started.
- Service failure: Read and write jobs on the lost device partition fail to run or run slowly.
- Service interruption: Customers may need time to restore data and systems, and services cannot be provided.
- Security risk: Important data may be stolen or disclosed, which severely affects customer services.

Possible Causes

- The hard disk is removed.
- The hard disk is offline, or a bad sector exists on the hard disk.

Procedure

- Step 1** On FusionInsight Manager, click **O&M > Alarm > Alarms**, and click  in the row where the alarm is located.
- Step 2** Obtain **HostName**, **PartitionName** and **DirName** from **Location**.
- Step 3** Check whether the disk of **PartitionName** on **HostName** is inserted to the correct server slot.
 - If yes, go to [Step 4](#).
 - If no, go to [Step 5](#).

- Step 4** Contact hardware engineers to remove the faulty disk.
- Step 5** Log in to the **HostName** node where an alarm is reported and check whether there is a line containing **DirName** in the **/etc/fstab** file as user **root**.
- If yes, go to [Step 6](#).
 - If no, go to [Step 7](#).
- Step 6** Run the **vi /etc/fstab** command to edit the file and delete the line containing **DirName**.
- Step 7** Contact hardware engineers to insert a new disk. For details, see the hardware product document of the relevant model. If the faulty disk is in a RAID group, configure the RAID group. For details, see the configuration methods of the relevant RAID controller card.
- Step 8** Wait 20 to 30 minutes (The disk size determines the waiting time), and run the **mount** command to check whether the disk has been mounted to the **DirName** directory.
- If yes, go to [Step 9](#) for MRS 3.3.0 and later, MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions. For clusters earlier than MRS 3.3.0, manually clear the alarm. No further action is required.
 - If no, go to [Step 10](#).
- Step 9** Wait about 2 minute and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 10](#).
- Collect fault information.**
- Step 10** On the FusionInsight Manager, choose **O&M > Log > Download**.
- Step 11** Select the **OmmServer** from the Services drop-down list and click **OK**.
- Step 12** Set Start Date for log collection to 10 minutes ahead of the alarm generation time and End Date to 10 minutes behind the alarm generation time and click **Download**.
- Step 13** Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

MRS 3.3.0 and later, MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions: After the fault is rectified, the system automatically clears this alarm.

Versions earlier than MRS 3.3.0: After the fault is rectified, the system does not automatically clear this alarm, and you need to manually clear the alarm.

Related Information

None

7.12.10 ALM-12015 Partition Filesystem Readonly

Description

The system checks the partition status every 60 seconds. This alarm is generated when the system detects that a partition to which service directories are mounted enters the read-only mode (due to a bad sector or a faulty file system). The system checks the partition status periodically.

This alarm is cleared when the system detects that the partition to which service directories are mounted exits from the read-only mode (because the file system is restored to read/write mode, the device is removed, or the device is formatted).

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12015 | Major | Yes |

Parameters

| Name | Meaning |
|---------------|---|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| DirName | Specifies the directory for which the alarm is generated. |
| PartitionName | Specifies the device partition for which the alarm is generated. |


Impact on the System

- Service failure: If a job needs to modify the data on the read-only device partition, the job may fail to run.
- Latency: If some components need to synchronize data to the read-only device partition, data synchronization may fail or time out, causing service delay.

Possible Causes

The hard disk is faulty, for example, a bad sector exists.

Procedure

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**, click  in the row where the alarm is located.
- Step 2** Obtain **HostName** and **PartitionName** from **Location**. **HostName** is the node where the alarm is reported, and **PartitionName** is the partition of the faulty disk.
- Step 3** Contact hardware engineers to check whether the disk is faulty. If the disk is faulty, remove it from the server.
- Step 4** After the disk is removed, alarm **ALM-12014 Partition Lost** is reported. Handle the alarm. For details, see [ALM-12014 Partition Lost](#). After the alarm **ALM-12014 Partition Lost** is cleared, alarm **ALM-12015 Partition Filesystem Readonly** is automatically cleared.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.11 ALM-12016 CPU Usage Exceeds the Threshold

Description

The system checks the CPU usage every 30 seconds and compares the actual CPU usage with the threshold. The CPU usage has a default threshold. This alarm is generated when the CPU usage exceeds the threshold for several times (configurable, 10 times by default) consecutively.

The alarm is cleared in the following two scenarios: The value of **Trigger Count** is 1 and the CPU usage is smaller than or equal to the threshold; the value of **Trigger Count** is greater than 1 and the CPU usage is smaller than or equal to 90% of the threshold.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12016 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

- Latency: If the CPU usage of a host is too high, service processes may run slowly and services may be delayed.
- Service failure: If the host CPU usage is too high, service processing may slow down, time out, or fail. As a result, jobs may fail to run.

Possible Causes

- The alarm threshold or alarm smoothing times are incorrect.
- CPU configuration cannot meet service requirements. The CPU usage reaches the upper limit. Or the service is in peak hours. As a result, the CPU usage reaches the upper limit in a short period of time.

Procedure

Check whether the alarm threshold or alarm Trigger Count are correct.


Step 1 Change the alarm threshold and alarm **Trigger Count** based on CPU usage.

On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > CPU > Host CPU Usage** and change the alarm smoothing times based on CPU usage, as shown in [Figure 7-52](#).

NOTE

This option defines the alarm check phase. **Trigger Count** indicates the alarm check threshold. An alarm is generated when the number of check times exceeds the threshold.

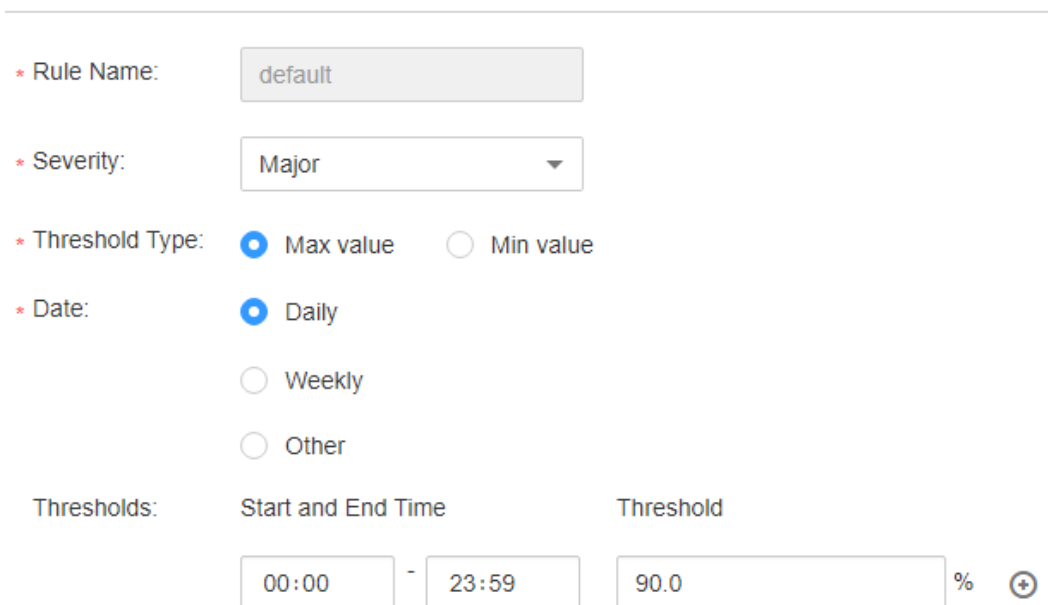
Figure 7-52 Setting alarm smoothing times



On **Host CPU Usage** page and click **Modify** in the **Operation** column to change the alarm threshold, as shown in [Figure 7-53](#).

Figure 7-53 Setting an alarm threshold

Thresholds > **Modify Rule**



Step 2 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

Check whether the CPU usage reaches the upper limit.

Step 3 In the alarm list on FusionInsight Manager, click in the row where the alarm is located to view the alarm host address in the alarm details.

Step 4 On the **Hosts** page, click the node on which the alarm is reported.

Step 5 View the CPU usage for 5 minutes. If the CPU usage exceeds the threshold for multiple times, contact the system administrator to add more CPUs.

Step 6 Check whether the current traffic is in peak hours. If the alarm is generated during peak hours, you are advised to expand the capacity of the node or contact the system administrator to add more CPUs.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

Collect fault information.

Step 8 On the FusionInsight Manager in the active cluster, choose **O&M > Log > Download**.

Step 9 Select **OmmServer** from the **Service** and click **OK**.

Step 10 Set **Start Date** for log collection to 10 minutes ahead of the alarm generation time and **End Date** to 10 minutes behind the alarm generation time in **Time Range** and click **Download**.

Step 11 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.12 ALM-12017 Insufficient Disk Capacity

Description

The system checks the host disk usage of the system every 30 seconds and compares the actual disk usage with the threshold. The disk usage has a default threshold, this alarm is generated when the host disk usage exceeds the specified threshold.

When the **Trigger Count** is 1, this alarm is cleared when the usage of a host disk partition is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the usage of a host disk partition is less than or equal to 90% of the threshold.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12017 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| PartitionName | Specifies the device partition for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

Service failure: If you need to modify or use data on the disk when the disk capacity is insufficient, the job may fail.

Possible Causes

- The alarm threshold is incorrect.
- Disk configuration of the server cannot meet service requirements.

Procedure

Check whether the alarm threshold is appropriate.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Disk > Disk Usage** and check whether the threshold (configurable, 90% by default) is appropriate.

- If yes, go to [Step 2](#).
- If no, go to [Step 4](#).

Step 2 Choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Disk > Disk Usage** and click **Modify** in the **Operation** column to change the alarm threshold based on site requirements. As shown in [Figure 7-54](#):

Figure 7-54 Setting an alarm threshold

Thresholds > **Modify Rule**

* Rule Name:

* Severity:

* Threshold Type: Max value Min value

* Date: Daily
 Weekly
 Other


Thresholds: Start and End Time Threshold

- %

Step 3 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Check whether the disk usage reaches the upper limit.

Step 4 In the alarm list on FusionInsight Manager, click  in the row where the alarm is located to view the alarm host name and disk partition information in the alarm details.

Step 5 Log in to the node where the alarm is generated as user **root**.

Step 6 Run the `df -lmPT | awk '$2 != "iso9660"' | grep '^/dev/' | awk '{"readlink -m "$1 | getline real }{$1=real; print $0}' | sort -u -k 1,1` command to check the system disk partition usage. Check whether the disk is mounted to the following directories based on the disk partition name obtained in [Step 4](#): `/`, `/opt`, `/tmp`, `/var`, `/var/log`, and `/srv/BigData` (can be customized).


- If yes, the disk is a system disk. Then go to [Step 10](#).
- If no, the disk is not a system disk. Then go to [Step 7](#).

Step 7 Run the `df -lmPT | awk '$2 != "iso9660"' | grep '^/dev/' | awk '{"readlink -m "$1 | getline real }{$1=real; print $0}' | sort -u -k 1,1` command to check the system disk partition usage. Determine the role of the disk based on the disk partition name obtained in [Step 4](#).

Step 8 Check the disk service.

In MRS, check whether the disk service is HDFS, Yarn, Kafka, Supervisor.

- If yes, adjust the capacity. Then go to [Step 9](#).
- If no, go to [Step 12](#).

- Step 9** After 2 minutes, check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 12](#).
- Step 10** Run the `find / -xdev -size +500M -exec ls -l {} \;` command to check whether a file larger than 500 MB exists on the node and disk.
- If yes, go to [Step 11](#).
 - If no, go to [Step 12](#).
- Step 11** Handle the large file and check whether the alarm is cleared 2 minutes later.
- If yes, no further action is required.
 - If no, go to [Step 12](#).
- Step 12** Contact the system administrator to expand the disk capacity.
- Step 13** After 2 minutes, check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 14](#).
- Collect fault information.**
- Step 14** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 15** Select **OMS** from the **Service** and click **OK**.
- Step 16** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 17** Contact the O&M personnel and send the collected log information.
- End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.13 ALM-12018 Memory Usage Exceeds the Threshold

Description

The system checks the memory usage of the system every 30 seconds and compares the actual memory usage with the threshold. The memory usage has a default threshold, this alarm is generated when the value of the memory usage exceeds the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the host memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the host memory usage is less than or equal to 90% of the threshold.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12018 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System


- Latency: If the host memory usage is too high, service processes may run slowly and services may be delayed.
- Service failure: If the host memory usage is too high, memory overflow may occur in service processes and jobs may fail.

Possible Causes

Memory configuration cannot meet service requirements. The memory usage reaches the upper limit.

Procedure

Expand the system.

- Step 1** In the alarm list on FusionInsight Manager, click  in the row where the alarm is located to view the alarm host address in the alarm details.
- Step 2** Log in to the host where the alarm is generated as user **root**.
- Step 3** If the memory usage exceeds the threshold, perform memory capacity expansion.
- Step 4** Run the command **free -m | grep Mem\:** | **awk '{printf("%s,", \$3 * 100 / \$2)}'** to check the system memory usage.


Step 5 Wait for 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager in the active cluster, choose **O&M > Log > Download**.

Step 7 Select **OmmServer** from the **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.14 ALM-12027 Host PID Usage Exceeds the Threshold

Description

The system checks the PID usage every 30 seconds and compares the actual PID usage with the default PID usage threshold. This alarm is generated when the system detects that the PID usage exceeds the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the PID usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the PID usage is less than or equal to 90% of the threshold.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12027 | Major | Yes |

Parameters

| Name | Meaning |
|--------|---|
| Source | Specifies the cluster or system for which the alarm is generated. |

| Name | Meaning |
|-------------------|--|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System


Service failure: If the PID usage of the host is too high, PIDs cannot be allocated to new service processes. As a result, jobs may fail to be executed.

Possible Causes

Too many processes are running on the node. You need to increase the value of `pid_max`.

Procedure

Increase the value of `pid_max`.

- Step 1** In the alarm list on FusionInsight Manager, click  in the row where the alarm is located to view the alarm host address in the alarm details.
- Step 2** Log in to the host where the alarm is generated as user **root**.
- Step 3** Run the `cat /proc/sys/kernel/pid_max` command to check the value of `pid_max`.
- Step 4** If the PID usage exceeds the threshold, edit the `/etc/sysctl.conf` file and increase the value of `kernel.pid_max` to twice of the value of `pid_max` queried in **Step 3**. If this parameter does not exist, add it to the end of the file.

For example, change the parameter to `kernel.pid_max=65536` and run the following command to make the parameter take effect immediately:

```
sysctl -p
```

NOTE

The maximum value of `kernel.pid_max` is as follows:

- On 32-bit systems: 32768
- On 64-bit systems: 4194304 (2²²)


- Step 5** Wait for 5 minutes, and check whether the alarm is cleared.
 - If yes, no further action is required.

- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager home page of the active cluster, choose **O&M > Log > Download**.

Step 7 Select all services from the **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.15 ALM-12028 Number of Processes in the D State and Z State on a Host Exceeds the Threshold

Alarm Description

The system checks the number of processes in the D state and Z state of user **omm** on the host every 30 seconds and compares the actual number with the threshold. The number of processes in the D state and Z state on the host has a default threshold range. This alarm is generated when the number of processes exceeds the threshold.

This alarm is cleared when **Trigger Count** is **1** and the total number of processes in the D state and Z state of user **omm** on the host does not exceed the threshold. This alarm is cleared when **Trigger Count** is greater than **1** and the total number of processes in the D state and Z state of user **omm** on the host is less than or equal to 90% of the threshold.

NOTE

The function of checking the number of processes in the Z state on the host applies to MRS 3.2.0 or later.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 12028 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------------|--|
| Source | Specifies the cluster or system for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System


- Latency: New service processes cannot be created. Concurrent task processing may be slow and services may be delayed.
- Service failure: New service processes cannot be created, which may cause job failures.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state.

Handling Procedure

Check the processes in the D state and Z state.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click  to view the IP address of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**. () Then run the **su - omm** command to switch to user **omm**.

Step 3 Run the following command as user **omm** to view the PID of the process that is in the D state and Z state:

```
ps -elf | grep -v "\[thread_checkio\]" | awk 'NR!=1 {print $2, $3, $4}' | grep omm | awk -F ' ' '{print $1, $3}' | grep -E "Z|D" | awk '{print $2}'
```

Step 4 Check whether the command output is empty.

- If yes, the service process is running properly. Then go to [Step 6](#).
- If no, go to [Step 5](#).

Step 5 Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)


Step 6 Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 8 Select **OMS** for **Service** and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.16 ALM-12033 Slow Disk Fault

Alarm Description

For MRS 3.3.0 and its later versions as well as MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions:

- For HDDs, the alarm is triggered when any of the following conditions is met:
 - By default, the system collects data every 3 seconds. The svctm latency reaches 1000 ms within 30 seconds in at least seven collection periods.
 - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 150 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
 - By default, the system collects data every 3 seconds. The svctm latency reaches 1000 ms within 30 seconds in at least seven collection periods.
 - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 20 ms within 300 seconds.

The collection period is 3 seconds, and the detection period is 30 or 300 seconds. This alarm is automatically cleared when neither of the preceding conditions is met for three consecutive detection periods (30 or 300 seconds).

For versions earlier than MRS 3.3.0:

- For HDDs, the alarm is triggered when any of the following conditions is met:

- The system runs the **iostat** command every 3 seconds by default, and detects that the **svctm** value exceeds 1000 ms in at least seven consecutive periods within 30 seconds.
- The system runs the **iostat** command every 3 seconds, and detects that more than 50% of I/Os take more than 150 ms within 300s.
- For SSDs, the alarm is triggered when any of the following conditions is met:
 - The system runs the **iostat** command every 3 seconds by default, and detects that the **svctm** value exceeds 1000 ms for at least 10 periods within 30 seconds.
 - The system runs the **iostat** command every 3 seconds by default, and detects that more than 60% of I/Os take more than 20 ms within 300s.

This alarm is automatically cleared when the preceding conditions have not been met for 15 minutes.

 **NOTE**

For details about how to obtain the **svctm** value, see [Related Information](#).

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|---|--------------|
| 12033 | <ul style="list-style-type: none">• Minor: MRS 3.3.0 and its later versions as well as MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions• Major: versions earlier than MRS 3.3.0 | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster or system for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| DiskName | Specifies the disk for which the alarm was generated. |

Impact on the System

- The system I/O performance deteriorates, which means slow response and low throughput. For example, job submission is slow, page responds slowly, interface response times out, and the system is in error or even crash.
- System fault: Customer services may be interrupted. The system may break down and the key information stored on the faulty disk may be lost.

Possible Causes

The disk is aged or has bad sectors.

Handling Procedure

Check the disk status.

- Step 1** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**.
- Step 2** View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is generated.
- Step 3** Check whether the node for which the alarm is generated is in a virtualization environment.
- If yes, go to [Step 4](#).
 - If no, go to [Step 7](#).
- Step 4** Check whether the storage performance provided by the virtualization environment meets the hardware requirements. Then, go to [Step 5](#).
- Step 5** Log in to the alarm node as user **root**, run the **df -h** command, and check whether the command output contains the value of the **DiskName** field.
- If yes, go to [Step 7](#).
 - If no, go to [Step 6](#).
- Step 6** Run the **lsblk** command to check whether the mapping between the value of **DiskName** and the disk has been created.

```
sda                8:0    0 27810G 0
├─sda1             8:1    0   509M 0 /boot
└─sda2             8:2    0 278.4G 0
   ├─system-opt (dm-0) 253:0   0   50G 0 /opt
   ├─system-root (dm-1) 253:1   0   50G 0 /
   ├─system-swap (dm-2) 253:2   0   50G 0
   └─system-var (dm-3) 253:3   0   50G 0 /var
```

- If yes, go to [Step 7](#).
 - If no, go to [Step 22](#).
- Step 7** Log in to the alarm node as user **root**, run the **lsscsi | grep "/dev/sd[x]"** command to view the disk information, and check whether RAID has been set up.

 NOTE

In the command, `/dev/sd[x]` indicates the disk name obtained in [Step 2](#).

Example:

```
lsscsi | grep "/dev/sda"
```

In the command output, if **ATA**, **SATA**, or **SAS** is displayed in the third line, the disk has not been organized into a RAID group. If other information is displayed, RAID has been set up.

- If yes, go to [Step 12](#).
- If no, go to [Step 8](#).

Step 8 Run the `smartctl -i /dev/sd[x]` command to check whether the hardware supports the SMART tool.

Example:

```
smartctl -i /dev/sda
```

In the command output, if "SMART support is: Enabled" is displayed, the hardware supports SMART. If "Device does not support SMART" or other information is displayed, the hardware does not support SMART.

- If yes, go to [Step 9](#).
- If no, go to [Step 16](#).

Step 9 Run the `smartctl -H --all /dev/sd[x]` command to check basic SMART information and determine whether the disk is working properly.

Example:

```
smartctl -H --all /dev/sda
```

Check the value of **SMART overall-health self-assessment test result** in the command output. If the value is **FAILED**, the disk is faulty and needs to be replaced. If the value is **PASSED**, check the value of **Reallocated_Sector_Ct** or **Elements in grown defect list**. If the value is greater than 100, the disk is faulty and needs to be replaced.

- If yes, go to [Step 10](#).
- If no, go to [Step 18](#).

Step 10 Run the `smartctl -l error -H /dev/sd[x]` command to check the Glist of the disk and determine whether the disk is normal.

Example:

```
smartctl -l error -H /dev/sda
```

Check the **Command/Feature_name** column in the command output. If **READ SECTOR(S)** or **WRITE SECTOR(S)** is displayed, the disk has bad sectors. If other errors occur, the disk circuit board is faulty. Both errors indicate that the disk is abnormal and needs to be replaced.

If "No Errors Logged" is displayed, no error log exists. You can trigger the disk SMART self-check.

- If yes, go to [Step 11](#).
- If no, go to [Step 18](#).

Step 11 Run the `smartctl -t long /dev/sd[x]` command to trigger the disk SMART self-check. After the command is executed, the time when the self-check is to be completed is displayed. After the self-check is completed, repeat [Step 9](#) and [Step 10](#) to check whether the disk is working properly.

Example:

```
smartctl -t long /dev/sda
```

- If yes, go to [Step 17](#).
- If no, go to [Step 18](#).

Step 12 Run the `smartctl -d [sat|scsi]+megaraid,[DID] -H --all /dev/sd[x]` command to check whether the hardware supports SMART.

 NOTE

- In the command, `[sat|scsi]` indicates the disk type. Both types need to be used.
- `[DID]` indicates the slot information. Slots 0 to 15 need to be used.

For example, run the following commands in sequence:

```
smartctl -d sat+megaraid,0 -H --all /dev/sda
```

```
smartctl -d sat+megaraid,1 -H --all /dev/sda
```

```
smartctl -d sat+megaraid,2 -H --all /dev/sda
```

...

Try the command combinations of different disk types and slot information. If "SMART support is: Enabled" is displayed in the command output, the disk supports SMART. Record the parameters of the disk type and slot information when a command is successfully executed. If "SMART support is: Enabled" is not displayed in the command output, the disk does not support SMART.

- If yes, go to [Step 13](#).
- If no, go to [Step 16](#).

Step 13 Run the `smartctl -d [sat|scsi]+megaraid,[DID] -H --all /dev/sd[x]` command recorded in [Step 12](#) to check basic SMART information and determine whether the disk is normal.

Example:

```
smartctl -d sat+megaraid,2 -H --all /dev/sda
```

Check the value of **SMART overall-health self-assessment test result** in the command output. If the value is **FAILED**, the disk is faulty and needs to be replaced. If the value is **PASSED**, check the value of **Reallocated_Sector_Ct** or **Elements in grown defect list**. If the value is greater than 100, the disk is faulty and needs to be replaced.

- If yes, go to [Step 14](#).
- If no, go to [Step 18](#).

Step 14 Run the `smartctl -d [sat|scsi]+megaraid,[DID] -l error -H /dev/sd[x]` command to check the Glist of the disk and determine whether the hard disk is working properly.

Example:

```
smartctl -d sat+megaraid,2 -l error -H /dev/sda
```

Check the **Command/Feature_name** column in the command output. If **READ SECTOR(S)** or **WRITE SECTOR(S)** is displayed, the disk has bad sectors. If other errors occur, the disk circuit board is faulty. Both errors indicate that the disk is abnormal and needs to be replaced.

If "No Errors Logged" is displayed, no error log exists. You can trigger the disk SMART self-check.

- If yes, go to [Step 15](#).
- If no, go to [Step 18](#).

Step 15 Run the `smartctl -d [sat|scsi]+megaraid,[DID] -t long /dev/sd[x]` command to trigger the disk SMART self-check. After the command is executed, the time when the self-check is to be completed is displayed. After the self-check is completed, repeat [Step 13](#) and [Step 14](#) to check whether the disk is working properly.

Example:

```
smartctl -d sat+megaraid,2 -t long /dev/sda
```

- If yes, go to [Step 17](#).
- If no, go to [Step 18](#).

Step 16 If the configured RAID controller card does not support SMART, the disk does not support SMART. In this case, use the check tool provided by the corresponding RAID controller card vendor to rectify the fault. Then go to [Step 17](#).

For example, LSI is a MegaCLI tool.

Step 17 On FusionInsight Manager, choose **O&M > Alarm > Alarms**, click **Clear** in the **Operation** column of the alarm, and check whether the alarm is reported on the same disk again.

If the alarm is reported for three times, replace the disk.

- If yes, go to [Step 18](#).
- If no, no further action is required.

Replace the disk.

Step 18 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**.

Step 19 View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is reported.

Step 20 Replace the disk.


Step 21 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 22](#).

Collect the fault information.

Step 22 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 23 Select **OMS** for **Service** and click **OK**.

Step 24 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 25 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

The **svctm** value can be obtained as follows:

- MRS 3.1.0:

Run the **iostat -x -t** command in the OS.

```
lomm@node-master1hxyk ~]$ iostat -x -t
Linux 3.10.0-862.14.1.5.el7.x86_64 (node-master1hxyk) 11/11/2022 _x86_64_ (4 CPU)

11/11/2022 03:35:20 PM
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           27.66    0.00   15.66    0.63    0.00   56.06

Device:            rrqm/s   wrqm/s     r/s     w/s    kB/s    kB/s  avgrq-sz  avgqu-sz   await  r_await  w_await   svctm   %util
vda                 0.13    29.26    1.71    23.51  187.56  608.08   63.11    0.91    36.02   50.86   34.94    0.64    1.62
vdb                 0.00    14.45    0.08    27.34    1.35   301.81   22.12    0.08    2.81   26.57    2.74    0.53    1.45
```

- Versions later than MRS 3.1.0:

$svctm = (tot_ticks_new - tot_ticks_old) / (rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old)$

- Versions earlier than MRS 3.3.0: If $rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0$, then **svctm = 0**.
- MRS 3.3.0 and its later versions as well as MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions:

When the detection period is 30 seconds, if $rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0$, then **svctm = 0**.

When the detection period is 300 seconds and $rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0$, if $tot_ticks_new - tot_ticks_old = 0$, then **svctm = 0**; otherwise, the value of **svctm** is infinite.

The parameters can be obtained as follows:

The system runs the **cat /proc/diskstats** command every 3 seconds to collect data. For example:

```

comm@ 0 vda 1101553 35446 83439787 3338546 28744856 48314024 1054257652 52667332 0 19569526 40342913 0 0 0 0
253 1 vda1 25494 54533791 2565698 3440004 6749340 215777628 12114542 0 6473005 11339691 0 0 0 0
253 2 vda2 15 0 108 136 0 0 0 0 150 129 0 0 0 0
253 5 vda5 22373 1364 2525122 79502 4212374 4104759 161597984 8145606 0 3598808 6239095 0 0 0 0
253 6 vda6 11145 314 529002 85050 259201 70368 4412408 321454 0 189336 259725 0 0 0 0
253 7 vda7 157987 105 3477434 149542 6507077 1028968 140666992 14349866 0 1679035 11116587 0 0 0 0
253 8 vda8 312935 8169 22366722 458354 12179958 34360589 531802640 17724858 0 9060731 11385470 0 0 0 0
253 16 vdb 275920 21939 15977738 2171665 39472291 28236575 2653825040 482230505 0 30580346 465962048 0 0 0 0
253 17 vdb1 275439 21939 15948866 2171472 31290400 28236555 2653824832 481837775 0 30036724 465855080 0 0 0 0
7 0 loop0 356 0 17442 150 0 0 0 0 149 105 0 0 0 0
comm@ 0 vda 1101553 35446 83439787 3338546 28747977 48319338 1054352084 52672715 0 19571460 40346640 0 0 0 0
253 1 vda1 25494 54533791 2565698 3440010 6750402 215791076 12115169 0 6474429 11339985 0 0 0 0
253 2 vda2 15 0 108 136 0 0 0 0 150 129 0 0 0 0
253 5 vda5 22373 1364 2525122 79502 4212822 4105244 161614088 8146153 0 3599216 6239432 0 0 0 0
253 6 vda6 11145 314 529002 85050 259245 70433 4413368 321489 0 189389 259730 0 0 0 0
253 7 vda7 157987 105 3477434 149542 6507759 1029060 140677872 14351373 0 1670157 11117724 0 0 0 0
253 8 vda8 312935 8169 22366722 458354 12181277 34364199 531855680 17727525 0 9061647 11387424 0 0 0 0
253 16 vdb 275920 21939 15977738 2171665 39477604 28238831 2653881640 482234435 0 30581946 465964144 0 0 0 0
253 17 vdb1 275439 21939 15948866 2171472 31293358 28238811 2653881432 481841639 0 30038274 465857164 0 0 0 0
7 0 loop0 356 0 17442 150 0 0 0 0 149 105 0 0 0 0
    
```

In the two commands:

In the data collected for the first time, the number in the fourth column is the **rd_ios_old** value, the number in the eighth column is the **wr_ios_old** value, and the number in the thirteenth column is the **tot_ticks_old** value.

In the data collected for the second time, the number in the fourth column is the **rd_ios_new** value, the number in the eighth column is the **wr_ios_new** value, and the number in the thirteenth column is the **tot_ticks_new** value.

In this case, the value of **svctm** is as follows:

$$(19571460 - 19569526) / ((1101553 + 28747977 - 1101553 - 28744856) = 0.6197$$

7.12.17 ALM-12034 Periodical Backup Failure

Description

The system executes the periodic backup task every 60 minutes. This alarm is generated when a periodical backup task fails to be executed. This alarm is cleared when the next backup task is executed successfully.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12034 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |

| Name | Meaning |
|----------|--|
| HostName | Specifies the host for which the alarm is generated. |
| TaskName | Specifies the task. |

Impact on the System

The periodic backup task failed, resulting in no available backup packages during the time period when the backup failed. When a system exception occurs and you need to use the backup package to restore data, no backup package is available during the failure period. As a result, data during the failure period cannot be restored.


Possible Causes

The alarm cause depends on the task details. Handle the alarm according to the logs and alarm details.

Procedure

Check whether the disk space is sufficient.

Step 1 In the FusionInsight Manager portal, click **O&M > Alarm > Alarms**.

Step 2 In the alarm list, click  in the row where the alarm is located and obtain **TaskName** from **Location**.

Step 3 Log in to the active node of the cluster as the **root** user and check the backup and restoration logs in `/var/log/Bigdata/controller/backup/`.

```
cd /var/log/Bigdata/controller/backup/
```

```
vi Log file name
```

Check whether the log file contains information similar to the following:

```
Upload backup files to *** file failed, error info: ***
```

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

Step 4 On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**. Locate the backup task based on **TaskName**, click **Configure** in the **Operation** column, and check whether all configuration items are correctly configured.

- If yes, go to [Step 7](#).
- If no, modify the configuration, save the modification, and go to [Step 5](#).

Step 5 Choose **More > Back Up Now** to start the backup task and check whether the backup task is successfully executed.

- If yes, go to [Step 6](#).

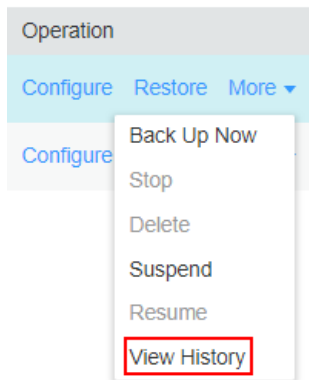
- If no, go to [Step 7](#).


Step 6 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Step 7 Choose **More > View History** in the **Operation** column. In the displayed dialog box view the task details.

Figure 7-55 View History



Step 8 In the displayed dialog box and click  to check whether the following message is displayed: Failed to backup xx due to insufficient disk space, move the data in the xx directory to other directories.

- If yes, go to [Step 9](#).
- If no, go to [Step 16](#).

Step 9 Choose **Backup Path > View** and obtain the **Backup Path**.

Step 10 Log in to the node as user **root** and run the following command to check the node mounting details:

df -h

Step 11 Check whether the available space of the node to which the backup path is mounted is less than 20 GB.

- If yes, go to [9](#).
- If no, go to [Step 16](#).

Step 12 Check whether there are many backup packages in the backup directory.

- If yes, go to [Step 13](#).
- If no, go to [Step 16](#).

Step 13 Enable the available space of the node to which the backup directory is mounted to be greater than 20 GB by moving backup packages out of the backup directory or delete the backup packages.

Step 14 After the problem is resolved, perform the backup task again and check whether the backup task execution is successful.

- If yes, go to [Step 15](#).
- If no, go to [Step 16](#).


Step 15 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 16](#).

Collect fault information.

Step 16 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 17 Select **Controller** from the **Service** and click **OK**.

Step 18 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.18 ALM-12035 Unknown Data Status After Recovery Task Failure

Description

After the recovery task fails, the system automatically rolls back every 60 minutes. If the rollback fails, data may be lost. If this occurs, an alarm is reported. This alarm is cleared when the next recovery task execution is successful.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12035 | Critical | Yes |

Parameters

| Name | Meaning |
|--------|---|
| Source | Specifies the cluster or system for which the alarm is generated. |

| Name | Meaning |
|-------------|---|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| TaskName | Specifies the task. |

Impact on the System


After the recovery task fails, the system automatically rolls back. If the rollback fails, data may be lost or the data status may be unknown, which may affect services.

Possible Causes

The alarm cause depends on the task details. Handle the alarm according to the logs and alarm details.

Procedure

Collect fault information.


- Step 1** In the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the running status of the component meets the requirements. (The OMS and DBService must be in the normal state, and other components must be stopped.)
- If yes, go to **Step 9**.
 - If no, go to **Step 2**.
- Step 2** Restore the component status as required and start the recovery task again.
- Step 3** Log in to the FusionInsight Manager portal and click **O&M** > **Alarm** > **Alarms**.
- Step 4** In the alarm list, click  in the row where the alarm is located to obtain **TaskName** from **Location**.
- Step 5** Choose **O&M** > **Backup and Restoration** > **Restoration Management**.
- Step 6** Find the restoration task by **Task Name** and view the task details.
- Step 7** Perform the recovery task again and check whether the recovery task execution is successful.
- If yes, go to **8**.
 - If no, go to **9**.
- Step 8** After 2 minutes, check whether the alarm is cleared.
- If yes, no further action is required.

- If no, go to 9.

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 10 Select **Controller** from the **Service** and click **OK**.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.19 ALM-12037 NTP Server Abnormal

Description

The system checks the NTP server status every 60 seconds. This alarm is generated when the system detects that the NTP server is abnormal for 10 consecutive times.

This alarm is cleared when the NTP server recovers.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12037 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |

| Name | Meaning |
|----------|--|
| HostName | Specifies the IP address of the NTP server for which the alarm is generated. |

Impact on the System


- The NTP server configured on the active OMS node is abnormal. In this case, the active OMS node cannot synchronize time with the NTP server and a time offset may be generated in the cluster.
- If the time difference exceeds 5 minutes, the client outside the cluster fails to authenticate the client inside the cluster. As a result, the job may fail to run.

Possible Causes

- The NTP server network is abnormal.
- The NTP server authentication fails.
- The NTP server time cannot be obtained.
- The time obtained from the NTP server is not continuously updated.

Procedure

Check the NTP server network.

- Step 1** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and click  in the row where the alarm is located.
- Step 2** View the alarm additional information to check whether the NTP server fails to be pinged.
- If yes, go to [Step 3](#).
 - If no, go to [Step 4](#).
- Step 3** Contact the network administrator to check the network configuration and ensure that the network between the NTP server and the active OMS node is normal. Then, check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 4](#).

Check whether the NTP server authentication fails.

- Step 4** Log in to the active OMS node as user **root**.
- Step 5** Run the following command to check the status of the resources on the active and standby nodes:
- ```
su - omm
sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh
```
- If "chrony" is displayed in the **ResName** column of the command output, go to [Step 6](#).

- If "ntp" is displayed in the **ResName** column, go to [Step 7](#).

 **NOTE**

If both "chrony" and "ntp" are displayed in the **ResName** column of the command output, the NTP service mode is being switched. Wait for 10 minutes and perform [Step 5](#) again. If both "chrony" and "ntp" still exist in the **ResName** column, contact O&M personnel.

**Step 6** Run the command **chronyc sources** to check whether the NTP server authentication fails.

If the value of **Reach** for chrony is **0**, the connection or authentication fails.

- If yes, go to [Step 12](#).
- If no, go to [Step 8](#).

**Step 7** Run the command **ntpq -np** to check whether the NTP server authentication fails.

If **refid** of the NTP server is **.AUTH.**, the authentication fails.

- If yes, go to [Step 12](#).
- If no, go to [Step 8](#).

**Check whether the time can be obtained from the NTP server.**

**Step 8** View the alarm additional information to check whether the time can be obtained from the NTP server.

- If yes, go to [Step 9](#).
- If no, go to [Step 10](#).

**Step 9** Contact the provider of the NTP server to rectify the NTP server fault. After the NTP server is normal, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

**Check whether the time obtained from the NTP server is not continuously updated.**

**Step 10** View the alarm additional information to check whether the time obtained from the NTP server is not continuously updated.

- If yes, go to [Step 11](#).
- If no, go to [Step 12](#).


**Step 11** Contact the provider of the NTP server to rectify the NTP server fault. After the NTP server is normal, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

**Collect fault information.**

**Step 12** On the FusionInsight Manager, choose **O&M > Log > Download**.

**Step 13** Select **NodeAgent** and **OmmServer** from the **Service** and click **OK**.

**Step 14** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 15** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.20 ALM-12038 Monitoring Indicator Dumping Failure

### Description

After monitoring indicator dumping is configured on FusionInsight Manager, the system checks the monitoring indicator dumping result at the dumping interval (60 seconds by default). This alarm is generated when the dumping fails.

This alarm is cleared when dumping is successful.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12038	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

The upper-layer management system cannot obtain monitoring indicators from the FusionInsight Manager system.



## Possible Causes

- The server cannot be connected.
- The save path on the server cannot be accessed.
- The monitoring indicator file fails to be uploaded.

## Procedure

### Check whether the server connection is normal.

**Step 1** Check whether the network between the FusionInsight Manager system and the server is normal.

- If yes, go to [Step 3](#).
- If no, go to [Step 2](#).

**Step 2** Contact the network administrator to recover the network and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

**Step 3** Choose **System > Interconnection > Upload Performance Data** and check whether the FTP username, password, port, dump mode, and public key configured on the upload performance data page are consistent with the configuration on the server.

- If yes, go to [Step 5](#).
- If no, go to [Step 4](#).

**Step 4** Enter the correct configuration information, click **OK**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

### Check the permission of the save path on the server is correct.

**Step 5** Choose **System > Interconnection > Upload Performance Data** and check the configuration items **FTP Username**, **Save Path**, and **Dump Mode**.

- If the dump mode is FTP, go to [Step 6](#).
- If the dump mode is SFTP, go to [Step 7](#).

**Step 6** Log in to the server in FTP mode. In the default path, check whether **FTP Username** has the read and write permission of the relative path **Save Path**.

- If yes, go to [Step 9](#).
- If no, go to [Step 8](#).

**Step 7** Log in to the server in SFTP mode and check whether **FTP Username** has the read and write permission of the absolute path **Save Path**.

- If yes, go to [Step 9](#).
- If no, go to [Step 8](#).

**Step 8** Add the read and write permission and check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to [Step 9](#).

**Check whether the save path on the server has sufficient disk space.**

**Step 9** Log in to the server and check whether the save path has sufficient disk space.

- If yes, go to [Step 11](#).
- If no, go to [Step 10](#).


**Step 10** Delete unnecessary files or go to the monitoring indicator dumping configuration page to change the save path. Then, check whether the save path has sufficient disk space.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Collect fault information.**

**Step 11** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 12** Select **OMS** from the **Service** and click **OK**.

**Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 14** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.21 ALM-12039 Active/Standby OMS Databases Not Synchronized

### Description

The system checks the data synchronization status between the active and standby OMS Databases every 10 seconds. This alarm is generated when the synchronization status cannot be queried for 30 consecutive times or when the synchronization status is abnormal.

This alarm is cleared when the data synchronization status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12039	Critical	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local GaussDB HA IP	Specifies the HA IP address of the local GaussDB.
Peer GaussDB HA IP	Specifies the HA IP address of the peer GaussDB.
SYNC_PERCENT	Specifies the synchronization percentage.

## Impact on the System


If the active/standby of the OMS database is not synchronized, data in the active database cannot be synchronized to the standby database. If the active instance is abnormal during the alarm reporting period, service data may be lost or data on the FusionInsight Manager may be abnormal.

## Possible Causes

- The network between the active and standby nodes is unstable.
- The standby OMS Database is abnormal.
- The standby node disk space is full.

## Procedure

**Check whether the network between the active and standby nodes is normal.**

- Step 1** Log in to FusionInsight Manager, click **O&M > Alarm > Alarms**, click  in the row where the alarm is located, and query the standby OMS Database IP address.
- Step 2** Log in to the active OMS Database node as user **root**.
- Step 3** Run the **ping Standby OMS Database heartbeat IP address** command to check whether the standby OMS Database node is reachable.
  - If yes, go to [Step 6](#).
  - If no, go to [Step 4](#).
- Step 4** Contact the network administrator to check whether the network is faulty.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** Rectify the network fault and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Check whether the standby OMS Database is normal.** (Skip this check for versions later than MRS 3.1.2.)

**Step 6** Log in to the standby OMS Database node as user **root**.

**Step 7** Run the **su - omm** command to switch to user **omm**.

**Step 8** Go to the `${BIGDATA_HOME}/om-server/om/sbin/` directory and run the `./status-oms.sh` command to check whether the OMS Database resource status of the standby DBService is normal. In the command output, check whether the following information is displayed in the row where **ResName** is **gaussDB**:

For example:

```
10_10_10_231 gaussDB Standby_normal Normal Active_standby
```

- If yes, go to [Step 9](#).
- If no, go to [Step 16](#).

**Check whether the standby node disk space is full.**

**Step 9** Log in to the standby OMS Database node as user **root**.

**Step 10** Run the **su - omm** command to switch to user **omm**.

**Step 11** Run the **echo \${BIGDATA\_DATA\_HOME}/dbdata\_om** command to obtain the OMS Database data directory.

**Step 12** Run the **df -h** command to view the system disk partition usage information.

**Step 13** Check whether the disk where the OMS Database data directory is mounted is full.

- If yes, go to [Step 14](#).
- If no, go to [Step 16](#).

**Step 14** Expand the disk capacity.


**Step 15** After the disk capacity is expanded, wait 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 16](#).

**Collect fault information.**

**Step 16** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 17** Select **OMMServer** from the **Service** and click **OK**.

**Step 18** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 19** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.22 ALM-12040 Insufficient System Entropy

### Alarm Description

MRS 3.2.0 or later:

The system checks whether the rng-tools or haveged tool has been enabled and correctly configured every 5 minutes. If neither tool is configured, this alarm is generated. If either is configured, the system continues to check the entropy. If the entropy is less than 100 for five consecutive times, this alarm is generated.

This alarm is cleared when rng-tools or haveged has been installed and enabled on the target node and the entropy of the OS is greater than or equal to 100 in at least one of five entropy checks.

MRS 3.1.2 or earlier:

The system checks the entropy for five consecutive times at 00:00 every day. Specifically, the system checks whether rng-tools or haveged has been enabled and correctly configured. If neither is configured, the system continues to check the entropy. If the entropy is less than 100 for five consecutive times, this alarm is reported.

This alarm is cleared when the system detects that the true random number mode has been configured, the random number parameters have been configured in the pseudo-random number mode, or neither mode is configured but the entropy of the OS is greater than or equal to 100 in at least one of five entropy checks.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12040	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

The entropy of the operating system of the node is insufficient. As a result, commands such as encryption and decryption are executed slowly on the node. As a result, the service processing performance of each instance deteriorates, and even service processes cannot be executed properly.

## Possible Causes

- rng-tools or haveged has not been installed or started.
- The entropy of the OS is smaller than 100 for multiple consecutive times.

## Handling Procedure

**Check whether haveged or rng-tools has been installed or started.**

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**.
- Step 2** Check the value of **HostName** in the **Location** area to obtain the name of the host for which the alarm is generated.
- Step 3** Log in to the node for which the alarm is generated as user **root**.
- Step 4** Run the **/bin/rpm -qa | grep -w "haveged"** command to check the haveged installation status and check whether the command output is empty.
- If yes, go to **Step 6**.
  - If no, go to **Step 5**.
- Step 5** Run the **/sbin/service haveged status |grep "running"** command and check the command output.
- If the command is executed successfully, haveged has been installed and configured correctly and is running properly. Go to **Step 8**.
  - If the command fails to execute, haveged is not running properly. Run the following command to manually restart haveged and go to **Step 9**:  
**systemctl restart haveged.service**
- Step 6** Run the **/bin/rpm -qa | grep -w "rng-tools"** command to check the rng-tools installation and check whether the command output is empty.
- If yes, contact the OS vendor to install and start haveged or rng-tools. Then go to **Step 9**.

- If no, go to [Step 7](#).

**Step 7** Run the `ps -ef | grep -v "grep" | grep rngd | tr -d " " | grep "\-r/dev/urandom"` command and check the command output.

- If the command is executed successfully, rngd has been installed and configured correctly and is running properly. Go to [Step 8](#).
- If the command fails to execute, rngd is not running properly. Run the following command to manually restart rngd and go to [Step 9](#):

```
systemctl restart rngd.service
```

#### Check the entropy of the OS.

**Step 8** Manually check the entropy of the OS.

Log in to the target node as user **root** and run the `cat /proc/sys/kernel/random/entropy_avail` command to check whether the entropy of the OS meets cluster installation requirements (no less than 100).

- If yes, the entropy of the OS is not less than 100. Go to [Step 9](#).
- If no, the entropy of the OS is less than 100. Use either of the following methods and go to [Step 9](#).
  - Method 1: Use haveged (true random number mode). Contact the OS vendor to install and start haveged.
  - Method 2: Use rng-tools (pseudo-random number mode). Contact the OS vendor to install and start rng-tools and configure it based on the OS type.


**Step 9** Wait until the system to check the entropy at 00:00 on the following day and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

#### Collect fault information.

**Step 10** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 11** Select **NodeAgent** for **Service** and click **OK**.

**Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## 7.12.23 ALM-12041 Incorrect Permission on Key Files

### Description

The system checks whether the permission, user, and user group information about critical directories or files is normal every 5 minutes. This alarm is generated when the information is abnormal.

This alarm is cleared when the information becomes normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12041	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
PathName	Specifies the path or name of the abnormal file.

### Impact on the System

System functions are unavailable.

- If the permission on the okerberos and oldap key files is abnormal, authentication fails and jobs may fail.
- If the permission on the controller and pms key files is abnormal, the process may be faulty, which may affect the elastic scaling performance.
- If the permission on key Tomcat files is abnormal, the login and viewing functions of FusionInsight Manager are affected.

### Possible Causes

The file permission is abnormal or the file is lost due to a user manually modified information such as the file permission, user, and user group, or the system is powered off unexpectedly.



## Procedure

Check whether the abnormal file exists and whether the permission on the abnormal file is correct.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**.
- Step 2** Check the value of **HostName** to obtain the host name involved in this alarm. Check the value of **PathName** to obtain the path or name of the abnormal file.
- Step 3** Log in to the node for which the alarm is generated as user **root**.
- Step 4** Run the **ll *pathName*** command, where *pathName* indicates the name of the abnormal file to obtain the user, permission, and user group information about the file or directory.
- Step 5** Go to **`\${BIGDATA\_HOME}/om-agent/nodeagent/etc/agent/autocheck** directory. Then run the **vi *keyfile*** command and search for the name of the abnormal file and check the due permission of the file.

### NOTE

To ensure proper configuration synchronization between the active and standby OMS servers, files, directories, and files and sub-directories in the directories configured in **`\${SOMS\_RUN\_PATH}/workspace/ha/module/hasync/plugin/conf/filesync.xml** will also be monitored except files and directories in **keyfile**. User **omm** must have read and write permissions of files and read and execute permissions of directories.

- Step 6** Compare the real-world permission of the file with the due permission obtained in [Step 5](#) and correct the permission, user, and user group information for the file.
- Step 7** Wait a hour and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 8](#).


### NOTE

If the disk partition where the cluster installation directory resides is used up, some temporary files will be generated in the program installation directory when running the **sed** command fails. Users do not have the read, write, and execute permissions of these temporary files. The system reports an alarm indicating that permissions of temporary files are abnormal if these files are within the monitoring range of the alarm. Perform the preceding alarm handling processes to clear the alarm. Alternatively, you can directly delete the temporary files after confirming that files with abnormal permissions are temporary. The temporary file generated after a **sed** command execution failure is similar to the following.

```
-rwx-----. 1 omm wheel 347 Jan 26 13:11 REALM_RESET_CONFIG
-rwx-----. 1 omm wheel 351 Jan 22 09:07 REALM_RESET_CONFIG_KRB
----- 1 omm wheel 0 Jan 26 13:15 sedbT8Cs4
-rwx-----. 1 omm wheel 7457 Jan 22 03:20 unlockuser.sh
```

Collect fault information.

- Step 8** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 9** Select **NodeAgent** from the **Service** and click **OK**.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.24 ALM-12042 Incorrect Configuration of Key Files

### Description

The system checks whether critical configurations are correct every 5 minutes. This alarm is generated when the configurations are abnormal.

This alarm is cleared when the configurations become normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12042	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
PathName	Specifies the path or name of the abnormal file.

## Impact on the System

Functions related to the file are abnormal.

- If the permission on the okerberos and oldap key files is abnormal, authentication fails and jobs may fail.
- If the permission on the controller and pms key files is abnormal, the process may be faulty, which may affect the elastic scaling performance.
- If the permission on key Tomcat files is abnormal, the login and viewing functions of FusionInsight Manager are affected.

## Possible Causes

The file configuration is modified manually or the system is powered off unexpectedly.

## Procedure

### Check abnormal file configuration.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**.
- Step 2** Check the value of **HostName** to obtain the host name involved in this alarm. Check the value of **PathName** to obtain the path or name of the abnormal file.
- Step 3** Log in to the node for which the alarm is generated as user **root**.
- Step 4** View the `$BIGDATA_LOG_HOME/nodeagent/scriptlog/checkfileconfig.log` file and analyze the cause based on the error log. Locate the check standards of the file in the [Related Information](#) and manually check and modify the file based on the standards.

Run the `vi file name` command to enter the editing mode, and then press **Insert** to start editing.


After the modification is complete, press **Esc** to exit the editing mode and enter `:wq` to save the settings and exit.

For example:

```
vi /etc/ssh/sshd_config
```

- Step 5** Wait a hour and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).

### Collect fault information.

- Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 7** Select **NodeAgent** from the **Service** and click **OK**.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

- **Check standards of /etc/fstab**  
Check whether the partitions configured in the **/etc/fstab** file can be found in **/proc/mounts**.  
Check whether the swap partitions configured in **fstab** correspond to those in **/proc/swaps**.
- **Check the /etc/hosts configuration file.**  
Run **cat /etc/hosts**. If any of the following situations occurs, the **/etc/hosts** configuration file is abnormal:
  - a. The **/etc/hosts** file does not exist.
  - b. The host name is not configured in the file.
  - c. The host name maps to multiple IP addresses in the file.
  - d. The IP address corresponding to the host name does not exist in the command output of the **ifconfig** command.
  - e. One IP address maps to multiple host names in the file.
- **Check standards of /etc/ssh/sshd\_config**  
Run the **vi /etc/ssh/sshd\_config** command to check whether configuration items are configured as follows:
  - a. The value of **UseDNS** must be set to **no**.
  - b. The value of **MaxStartups** must be greater than or equal to 1000.
  - c. At least one of the **PasswordAuthentication** and **ChallengeResponseAuthentication** parameters must be left blank or at least one of the parameters be set to **yes**.

## 7.12.25 ALM-12045 Read Packet Dropped Rate Exceeds the Threshold

### Alarm Description

The system checks the read packet dropped rate every 30 seconds. This alarm is generated when the read packet dropped rate exceeds the threshold (the default threshold is 0.5%) for multiple times (the default value is 5).

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Packet Dropped Rate**.

This alarm is cleared when **Trigger Count** is 1 and the read packet dropped rate is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the read packet dropped rate is less than or equal to 90% of the threshold.

The alarm detection is disabled by default. If you want to enable this function, check whether this function can be enabled based on Checking System Environments.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12045	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
PortName	Specifies the network port for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System


- Latency: When the read packet loss rate of the host network exceeds the threshold, request response is slowed down and services are delayed.
- Service failure: When the read packet loss rate of the host network exceeds the threshold, requests cannot be properly responded or times out, which may cause job running failures.

Risk warning: In SUSE kernel 3.0 or later or Red Hat 7.2, the system kernel modifies the mechanism for counting the number of dropped read packets. In this case, this alarm may be generated even if the network is running properly, but services are not affected. You are advised to check the system environment first.

## Possible Causes

- The NICs are bonded in active/standby mode.
- The alarm threshold is improperly configured.
- The network quality is poor.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the page that is displayed, click  in the row containing the alarm, and view the name of the host for which the alarm is generated and the NIC name.

**Check whether the NICs are bonded in active/standby mode.**

**Step 2** Log in to the alarm node as user **omm** and run the **ls -l /proc/net/bonding** command to check whether the **/proc/net/bonding** directory exists on the node.

- If yes, the bond mode is configured for the node. Go to [Step 3](#).

```
ls -l /proc/net/bonding/
total 0
-r--r--r-- 1 root root 0 Oct 11 17:35 bond0
```

- If no, the bond mode is not configured for the node. Go to [Step 5](#).

```
ls -l /proc/net/bonding/
ls: cannot access /proc/net/bonding/: No such file or directory
```

**Step 3** Run the **cat /proc/net/bonding/bond0** command to check whether the value of **Bonding Mode** in the configuration file is **fault-tolerance**.

### NOTE

In the command, **bond0** indicates the name of the bond configuration file. Use the file name obtained in [Step 2](#).

```
cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
```

```
Bonding Mode: fault-tolerance (active-backup)
Primary Slave: eth1 (primary_reselect always)
Currently Active Slave: eth1
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
```

```
Slave Interface: eth0
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 1
Slave queue ID: 0
```

```
Slave Interface: eth1
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 1
Slave queue ID: 0
```

- If yes, the NICs are bonded in active/standby mode. Go to [Step 4](#).
- If no, go to [Step 5](#).

**Step 4** Check whether the NIC specified by **NetworkCardName** in the alarm is the standby NIC.

- If yes, the alarm of the standby NIC cannot be automatically cleared. Manually clear the alarm on the alarm management page. No further action is required.
- If no, go to [Step 5](#).

**NOTE**

To determine the standby NIC, check the `/proc/net/bonding/bond0` configuration file. If the NIC name corresponding to `NetworkCardName` is **Slave Interface** but not **Currently Active Slave** (the current active NIC), the NIC is the standby one.

**Check whether the threshold is set properly.**

**Step 5** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Packet Dropped Rate**, and check whether the alarm threshold is configured properly. The default value is **0.5%**. You can adjust the threshold as needed.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

**Step 6** Choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Packet Dropped Rate**. Click **Modify** in the **Operation** column to change the threshold. See **Figure 7-56**.

**Figure 7-56** Configuring the alarm threshold

Thresholds > **Modify Rule**

---

\* Rule Name:

\* Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Other

Thresholds:	Start and End Time	Threshold
	<input type="text" value="00:00"/> - <input type="text" value="23:59"/>	<input type="text" value="0.5"/> % <input type="button" value="⊕"/>

**Step 7** After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

**Check whether the network connection is normal.**

**Step 8** Contact the network administrator to check whether the network is normal.

- If yes, rectify the fault and go to **Step 9**.
- If no, go to **Step 10**.

**Step 9** After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.


- If no, go to [Step 10](#).

**Collect the fault information.**

**Step 10** On FusionInsight Manager of the active cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 11** Select **OMS** for **Service** and click **OK**.

**Step 12** Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

**Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 14** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.26 ALM-12046 Write Packet Dropped Rate Exceeds the Threshold

### Alarm Description

The system checks the write packet dropped rate every 30 seconds. This alarm is generated when the write packet dropped rate exceeds the threshold (the default threshold is 0.5%) for multiple times (the default value is 5).

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Packet Dropped Rate**.

If **Trigger Count** is **1**, this alarm is cleared when the network write packet dropped rate is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the network write packet dropped rate is less than or equal to 90% of the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12046	Major	Yes



## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Port Name	Specifies the network port for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

- Latency: Requests are responded slowly and services are delayed.
- Service failure: Requests cannot be responded or time out. Jobs may fail to run.

## Possible Causes

- The alarm threshold is improperly configured.
- The network quality is poor.

## Handling Procedure

**Check whether the threshold is set properly.**

**Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Packet Dropped Rate**, and check whether the alarm threshold is configured properly. The default value is **0.5%**. You can adjust the threshold as needed.

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

**Step 2** Choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Packet Dropped Rate**. Click **Modify** in the **Operation** column to change the threshold.

See [Figure 7-57](#).

**Figure 7-57** Configuring the alarm threshold

Thresholds > **Modify Rule**

---


\* Rule Name:

\* Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Other

Thresholds: Start and End Time      Threshold

-        % 

**Step 3** After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the network connection is normal.**

**Step 4** Contact the network administrator to check whether the network is normal.

- If yes, rectify the fault and go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** After 5 minutes, check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect the fault information.**

**Step 6** On FusionInsight Manager of the active cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Select **OMS** for **Service** and click **OK**.

**Step 8** Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 10** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.27 ALM-12047 Read Packet Error Rate Exceeds the Threshold

### Alarm Description

The system checks the read packet error rate every 30 seconds. This alarm is generated when the read packet error rate exceeds the threshold (the default threshold is **0.5%**) for multiple times (the default value is **5**).

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Packet Error Rate**.

If **Trigger Count** is **1**, this alarm is cleared when the read packet error rate is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the read packet error rate is less than or equal to 90% of the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12047	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Port Name	Specifies the network port for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

- Latency: Requests are responded slowly and services are delayed.
- Service failure: Requests cannot be responded or time out. Jobs may fail to run.

## Possible Causes

- The alarm threshold is improperly configured.
- The network quality is poor.

## Handling Procedure

Check whether the threshold is set properly.

**Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Packet Error Rate**, and check whether the alarm threshold is configured properly. The default value is **0.5%**. You can adjust the threshold as needed.

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

**Step 2** Choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Packet Error Rate**. Click **Modify** in the **Operation** column to change the threshold.

See [Figure 7-58](#).

**Figure 7-58** Configuring the alarm threshold

Thresholds > **Modify Rule**

---


\* Rule Name:

\* Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Other

Thresholds: Start and End Time      Threshold

-        % 

**Step 3** After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the network connection is normal.**

**Step 4** Contact the network administrator to check whether the network is normal.

- If yes, rectify the fault and go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** After 5 minutes, check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect the fault information.**

**Step 6** On FusionInsight Manager of the active cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Select **OMS** for **Service** and click **OK**.

**Step 8** Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 10** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.28 ALM-12048 Write Packet Error Rate Exceeds the Threshold

### Alarm Description

The system checks the write packet error rate every 30 seconds. This alarm is generated when the write packet error rate exceeds the threshold (the default threshold is **0.5%**) for multiple times (the default value is **5**).

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Packet Error Rate**.

If **Trigger Count** is **1**, this alarm is cleared when the write packet error rate is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the write packet error rate is less than or equal to 90% of the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12048	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Port Name	Specifies the network port for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

- Latency: Requests are responded slowly and services are delayed.
- Service failure: Requests cannot be responded or time out. Jobs may fail to run.

## Possible Causes

- The alarm threshold is improperly configured.
- The network quality is poor.

## Handling Procedure

**Check whether the threshold is set properly.**

**Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Packet Error Rate**, and check whether the alarm threshold is configured properly. The default value is **0.5%**. You can adjust the threshold as needed.

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

**Step 2** Choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Packet Error Rate**. Click **Modify** in the **Operation** column to change the threshold.

See [Figure 7-59](#).

**Figure 7-59** Configuring the alarm threshold

Thresholds > Modify Rule

---


\* Rule Name:

\* Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Other

Thresholds: Start and End Time Threshold

-   % 

**Step 3** After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the network connection is normal.**

**Step 4** Contact the network administrator to check whether the network is normal.

- If yes, rectify the fault and go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** After 5 minutes, check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect the fault information.**

**Step 6** On FusionInsight Manager of the active cluster, choose **O&M > Log > Download**.

**Step 7** Select **OMS** for **Service** and click **OK**.

**Step 8** Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 10** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.29 ALM-12049 Network Read Throughput Rate Exceeds the Threshold

### Description

The system checks the network read throughput rate every 30 seconds and compares the actual throughput rate with the threshold (the default threshold is 80%). This alarm is generated when the system detects that the network read throughput rate exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Throughput Rate**.

When the **Trigger Count** is 1, this alarm is cleared when the network read throughput rate is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the network read throughput rate is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12049	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.



Name	Meaning
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

- Latency: When the host network read throughput exceeds the threshold, the request response slows down, causing service delay.
- Service failure: When the host network read throughput exceeds the threshold, requests cannot be properly responded or timed out, which may cause job execution failures.

## Possible Causes

- The alarm threshold is set improperly.
- The network port rate cannot meet the current service requirements.

## Procedure

### Check whether the threshold is set properly.

- Step 1** On the FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Throughput Rate** and check whether the alarm threshold is set properly. (By default, 80% is a proper value. However, users can configure the value as required.)
- If yes, go to [Step 2](#).
  - If no, go to [Step 4](#).
- Step 2** Based on actual usage condition, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Reading > Read Throughput Rate** and click **Modify** in the **Operation** column to modify the alarm threshold.

For details, see [Figure 7-60](#).

**Figure 7-60** Setting alarm thresholdsThresholds > **Modify Rule**

---


\* Rule Name:

\* Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Other


Thresholds:      Start and End Time      Threshold

-        % 

**Step 3** Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the network port rate can meet the service requirements.**

**Step 4** On FusionInsight Manager, click  in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host and the network port name for which the alarm is generated.

**Step 5** Log in to the host for which the alarm is generated as user **root**.

**Step 6** Run the **ethtool network port name** command to check the maximum speed of the current network port.

 **NOTE**

In the VM environment, you cannot run a command to query the network port rate. It is recommended that you contact the system administrator to confirm whether the network port rate meets the requirements.


**Step 7** If the network read throughput rate exceeds the threshold, contact the system administrator to increase the network port rate.

**Step 8** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Collect fault information.**

**Step 9** On the FusionInsight Manager home page of the active cluster, choose **O&M > Log > Download**.

- Step 10** Select **OMS** from the **Service** and click **OK**.
- Step 11** Set **Host** to the node for which the alarm is generated and the active OMS node.
- Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 13** Contact the O&M personnel and send the collected log information.
- End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.30 ALM-12050 Network Write Throughput Rate Exceeds the Threshold

### Description

The system checks the network write throughput rate every 30 seconds and compares the actual throughput rate with the threshold (the default threshold is 80%). This alarm is generated when the system detects that the network write throughput rate exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Throughput Rate**.

When the **Trigger Count** is 1, this alarm is cleared when the network write throughput rate is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the network write throughput rate is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12050	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

- Latency: When the host network write throughput exceeds the threshold, the request response is slowed down and services are delayed.
- Service failure: When the host network write throughput exceeds the threshold, requests cannot be responded or times out, which may cause job running failures.

## Possible Causes

- The alarm threshold is set improperly.
- The network port rate cannot meet the current service requirements.

## Procedure

Check whether the threshold is set properly.

**Step 1** On the FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Throughput Rate** and check whether the alarm threshold is set properly. (By default, 80% is a proper value. However, users can configure the value as required.)

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

**Step 2** Based on actual usage condition, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Network Writing > Write Throughput Rate** and click **Modify** in the **Operation** column to modify the alarm threshold.

For details, see [Figure 7-61](#).

**Figure 7-61** Setting alarm thresholdsThresholds > **Modify Rule**

---


\* Rule Name:

\* Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Other


Thresholds:      Start and End Time      Threshold

-        % 

**Step 3** Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the network port rate can meet the service requirements.**

**Step 4** On FusionInsight Manager, click  in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host and the network port name for which the alarm is generated.

**Step 5** Log in to the host for which the alarm is generated as user **root**.

**Step 6** Run the `ethtool network port name` command to check the maximum speed of the current network port.

 **NOTE**

In the VM environment, you cannot run a command to query the network port rate. It is recommended that you contact the system administrator to confirm whether the network port rate meets the requirements.


**Step 7** If the network write throughput rate exceeds the threshold, contact the system administrator to increase the network port rate.

**Step 8** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Collect fault information.**

**Step 9** On the FusionInsight Manager home page of the active cluster, choose **O&M > Log > Download**.

- Step 10** Select **OMS** from the **Service** and click **OK**.
- Step 11** Set **Host** to the node for which the alarm is generated and the active OMS node.
- Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 13** Contact the O&M personnel and send the collected log information.
- End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.31 ALM-12051 Disk Inode Usage Exceeds the Threshold

### Description

The system checks the disk Inode usage every 30 seconds and compares the actual Inode usage with the threshold (the default threshold is 80%). This alarm is generated when the Inode usage exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Disk > Disk Inode Usage**.

When the **Trigger Count** is 1, this alarm is cleared when the disk Inode usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the disk Inode usage is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12051	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System


Service failure: If you need to modify or use data on the disk when data cannot be written to the file system, the job may fail.

## Possible Causes

Massive small files are stored in the disk.

## Procedure

### Massive small files are stored in the disk.

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms** and click  in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host and the disk partition for which the alarm is generated.

**Step 2** Log in to the host for which the alarm is generated as user **root**.

**Step 3** Run the `df -i | grep -iE "partition name/Filesystem"` command to check the current disk Inode usage.

```
df -i | grep -iE "xvda2|Filesystem"
Filesystem Inodes IUsed IFree IUse% Mounted on
/dev/xvda2 2359296 207420 2151876 9% /
```

**Step 4** If the Inode usage exceeds the threshold, manually check small files stored in the disk partition and confirm whether these small files can be deleted.

### NOTE

Run the `for i in /*; do echo $i; find $i|wc -l; done` command to query the number of files in a partition. Replace `/*` with the specified partition.

```
for i in /srv/*; do echo $i; find $i|wc -l; done
/srv/BigData
4284
/srv/ftp
1
/srv/www
13
```

- If yes, run the **rm -rf *Path of the file or folder*** to be deleted command to delete the file or folder and go to [Step 5](#).

 **NOTE**

Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.

- If no, expand the capacity. Then, perform [Step 5](#).

**Step 5** Wait for 5 minutes, and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect fault information.**

**Step 6** On the FusionInsight Manager home page of the active cluster, choose **O&M > Log > Download**.

**Step 7** Select **OMS** from the **Service** and click **OK**.

**Step 8** Set **Host** to the node for which the alarm is generated and the active OMS node.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.32 ALM-12052 TCP Temporary Port Usage Exceeds the Threshold

### Description

The system checks the TCP temporary port usage every 30 seconds and compares the actual usage with the threshold (the default threshold is 80%). This alarm is generated when the TCP temporary port usage exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm > Thresholds > *Name of the desired cluster* > Host > Network Status > TCP Ephemeral Port Usage**.

When the **Trigger Count** is 1, this alarm is cleared when the TCP temporary port usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the TCP temporary port usage is less than or equal to 90% of the threshold.



## Attribute

Alarm ID	Alarm Severity	Auto Clear
12052	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System


Services on the host cannot establish external connections, and therefore they are interrupted.

## Possible Causes

- The temporary port cannot meet the current service requirements.
- The system is abnormal.

## Procedure

### Expand the temporary port number range.

- Step 1** On FusionInsight Manager, click  in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host for which the alarm is generated.
- Step 2** Log in to the host for which the alarm is generated as user **omm**.
- Step 3** Run the `cat /proc/sys/net/ipv4/ip_local_port_range |cut -f 1` command to obtain the value of the start port and run the `cat /proc/sys/net/ipv4/ip_local_port_range |cut -f 2` command to obtain the value of the end port. The total number of temporary ports is the value of the end port minus the value of the start port. If the total number of temporary ports is smaller than 28,232, the

random port range of the OS is narrow. Contact the system administrator to increase the port range.

**Step 4** Run the `ss -ant 2>/dev/null | grep -v LISTEN | awk 'NR > 2 {print $4}' | awk -F':' '{print $NF}' | awk '$1 >' "Value of the start port" {print $1}' | sort -u | wc -l` command to calculate the number of used temporary ports.

**Step 5** The formula for calculating the usage of the temporary ports is: Usage of the temporary ports = (Number of used temporary ports/Total number of temporary ports) x 100%. Check whether the temporary port usage exceeds the threshold.

- If yes, go to [Step 7](#).
- If no, go to [Step 6](#).

**Step 6** Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Check whether the system environment is abnormal.**

**Step 7** Run the following command to import the temporary file and view the frequently used ports in the `port_result.txt` file:

```
netstat -tnp|sort > $BIGDATA_HOME/tmp/port_result.txt
```

```
netstat -tnp|sort
Active Internet connections (w/o servers)

Proto Recv Send LocalAddress ForeignAddress State PID/ProgramName tcp 0 0 10-120-85-154:45433
10-120-85-154:9866 CLOSE_WAIT 94237/java
tcp 0 0 10-120-85-154:45434 10-120-85-154:9866 CLOSE_WAIT 94237/java
tcp 0 0 10-120-85-154:45435 10-120-85-154:9866 CLOSE_WAIT 94237/java
...
```

**Step 8** Run the following command to view the processes that occupy a large number of ports:

```
ps -ef |grep PID
```

#### NOTE

- PID is the processes ID queried in [Step 7](#).
- Run the following command to collect information about all processes and check the processes that occupy a large number of ports:

```
ps -ef > $BIGDATA_HOME/tmp/ps_result.txt
```


**Step 9** After obtaining the administrator's approval, clear the processes that occupy a large number of ports. Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

**Collect fault information.**

**Step 10** On the FusionInsight Manager home page of the active cluster, choose **O&M** > **Log** > **Download**.

**Step 11** Select **OMS** from the **Service** and click **OK**.

- Step 12** Set **Host** to the node for which the alarm is generated and the active OMS node.
- Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 14** Contact the O&M personnel and send the collected log information and files **port\_result.txt** and **ps\_result.txt**. Then, delete the two residual temporary files from the environment.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.33 ALM-12053 Host File Handle Usage Exceeds the Threshold

### Description

The system checks the file handle usage every 30 seconds and compares the actual usage with the threshold (the default threshold is 80%). This alarm is generated when the host file handle usage exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Host Status > Host File Handle Usage**.

When the **Trigger Count** is 1, this alarm is cleared when the host file handle usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the host file handle usage is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12053	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System


Service failure: When the host file handle usage exceeds the threshold, system applications cannot perform I/O operations such as file opening and network operations. As a result, the program is abnormal, which may cause job running failure.

## Possible Causes

- The application process is abnormal. For example, the opened file or socket is not closed.
- The number of file handles cannot meet the current service requirements.
- The system is abnormal.

## Procedure

### Check information about files opened in processes.

**Step 1** On FusionInsight Manager, click  in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host for which the alarm is generated.

**Step 2** Log in to the host for which the alarm is generated as user **root**.

**Step 3** Run the following command to check the process that occupies excessive file handles.

```
for proc in /proc/[0-9]*; do if [-d "$proc/fd"]; then num_fds=$(ls -l "$proc/fd" | wc -l); pid=$(basename $proc); echo "$num_fds ${pid}" ; fi; done | sort -nr | more
```

**Step 4** Check whether the processes in which a large number of files are opened are normal. For example, check whether there are files or sockets not closed.


- If yes, go to [Step 5](#).
- If no, go to [Step 7](#).

**Step 5** Release the abnormal processes that occupy too many file handles.

**Step 6** Five minutes later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Increase the number of file handles.**

**Step 7** On FusionInsight Manager, click  in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host for which the alarm is generated.

**Step 8** Log in to the host for which the alarm is generated as user **root**.

**Step 9** Contact the system administrator to increase the number of system file handles.

**Step 10** Run the `cat /proc/sys/fs/file-nr` command to view the used handles and the maximum number of file handles. The first value is the number of used handles, the third value is the maximum number. Please check whether the usage exceeds the threshold.

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

```
cat /proc/sys/fs/file-nr
12704 0 640000
```

**Step 11** Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

**Check whether the system environment is abnormal.**

**Step 12** Contact the system administrator to check whether the operating system is abnormal.

- If yes, go to [Step 13](#) to rectify the fault.
- If no, go to [Step 14](#).

**Step 13** Wait for 5 minutes, and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 14](#).

**Collect fault information.**

**Step 14** On the FusionInsight Manager home page of the active cluster, choose **O&M > Log > Download**.

**Step 15** Select **OMS** from the **Service** and click **OK**.

**Step 16** Set **Host** to the node for which the alarm is generated and the active OMS node.

**Step 17** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 18** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.34 ALM-12054 Invalid Certificate File

### Alarm Description

The system checks whether the certificate file is invalid (has expired or is not valid yet) on 23:00 every day. This alarm is generated when the certificate file is invalid.

This alarm is cleared when a valid certificate is imported and the alarm detection mechanism is triggered on the next hour.

#### NOTE

For MRS 3.2.0 or later, the certificate file is checked at the beginning of each hour.  
For versions earlier than MRS 3.2.0, the certificate file is checked on 23:00 every day.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12054	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System


The functions of the related modules cannot be used.

## Possible Causes

No certificate (CA certificate, HA root certificate, HA user certificate, Gaussdb root certificate, or Gaussdb user certificate) is imported to the system, the certificate fails to be imported, or the certificate file is invalid.

## Handling Procedure

**Check the alarm cause.**

**Step 1** On FusionInsight Manager, locate the target alarm in the real-time alarm list and click .

View **Additional Information** to obtain the additional information about the alarm.

- If **CA Certificate** is displayed in the additional alarm information, log in to the active OMS management node as user **omm** and go to [Step 2](#).
- If **HA root Certificate** is displayed in the additional information, view **Location** to obtain the name of the host involved in this alarm. Then, log in to the host as user **omm** and go to [Step 3](#).
- If **HA server Certificate** is displayed in the additional information, view **Location** to obtain the name of the host involved in this alarm. Then, log in to the host as user **omm** and go to [Step 4](#).
- If **Certificate has expired** is displayed in the additional information, view **Location** to obtain the name of the host for which the alarm is generated. Then, log in to the host as user **omm** and perform [Step 2](#) to [Step 4](#) in sequence to check whether the certificates have expired. If these certificates have not expired, check whether other certificates have been imported. If yes, import the certificate files again.

**Check the validity period of the certificate files in the system.**

**Step 2** Check whether the current system time is in the validity period of the CA certificate.

Run the **bash \${CONTROLLER\_HOME}/security/cert/conf/querycertvalidity.sh** command to check the effective time and due time of the CA root certificate.

- If yes, go to [Step 7](#).
- If no, go to [Step 5](#).

**Step 3** Check whether the current system time is in the validity period of the HA root certificate.

Run the **openssl x509 -noout -text -in \${CONTROLLER\_HOME}/security/certHA/root-ca.crt** command to check the effective time and due time of the HA root certificate.

- If yes, go to [Step 7](#).
- If no, go to [Step 6](#).

**Step 4** Check whether the current system time is in the validity period of the HA user certificate.

Run the **openssl x509 -noout -text -in \${CONTROLLER\_HOME}/security/certHA/server.crt** command to check the effective time and due time of the HA user certificate.

- If yes, go to [Step 7](#).
- If no, go to [Step 6](#).

The following is an example of the effective time and due time of a CA or HA certificate:

Certificate:

```
Data:
 Version: 3 (0x2)
 Serial Number:
 97:d5:0e:84:af:ec:34:d8
 Signature Algorithm: sha256WithRSAEncryption
 Issuer: C=CN, ST=xxx, L=yyy, O=zzz, OU=IT, CN=HADOOP.COM
 Validity
 Not Before: Dec 13 06:38:26 2016 GMT // Effective time
 Not After : Dec 11 06:38:26 2026 GMT // Due time
```

### Import certificate files.

**Step 5** Import a new CA certificate file.

Apply for or generate a new CA certificate file and import it to the system. For details, see [Replacing the CA Certificate](#). The alarm is automatically cleared after the CA certificate is imported. Check whether this alarm is reported again during periodic check.

- If yes, go to [Step 7](#).
- If no, no further action is required.

**Step 6** Import a new HA certificate file.


Apply for or generate a new HA certificate file and import it to the system. For details, see [Replacing the HA Certificate](#). The alarm is automatically cleared after the CA certificate is imported. Check whether this alarm is reported again during periodic check.

- If yes, go to [Step 7](#).
- If no, no further action is required.

### Collect the fault information.

**Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 8** In the **Services** area, select **Controller**, **OmmServer**, **OmmCore**, and **Tomcat**, and click **OK**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact O&M personnel and provide the collected logs.

----End



## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

For details about how to handle an expired OBS certificate, see [OBS Certificate in a Cluster Expired](#).

## 7.12.35 ALM-12055 Certificate File Is About to Expire

### Alarm Description

The system checks the certificate file on 23:00 every day. This alarm is generated if the certificate file is about to expire within 30 days.

This alarm is cleared when a certificate that is not about to expire is imported and the alarm detection mechanism is triggered on the next hour.

#### NOTE

For MRS 3.2.0 or later, the certificate file is checked at the beginning of each hour.  
For versions earlier than MRS 3.2.0, the certificate file is checked on 23:00 every day.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12055	Minor	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System


Some functions will be unavailable if the certificate expires.

## Possible Causes

The remaining validity period of a system certificate (CA certificate, HA root certificate, HA user certificate, Gaussdb root certificate, or Gaussdb user certificate) is less than 30 days.

## Handling Procedure

**Check the alarm cause.**

**Step 1** On FusionInsight Manager, locate the target alarm in the real-time alarm list and click .

View **Additional Information** to obtain the additional information about the alarm.

- If **CA Certificate** is displayed in the additional alarm information, log in to the active OMS management node as user **omm** and go to [Step 2](#).
- If **HA root Certificate** is displayed in the additional information, view **Location** to obtain the name of the host involved in this alarm. Then, log in to the host as user **omm** and go to [Step 3](#).
- If **HA server Certificate** is displayed in the additional information, view **Location** to obtain the name of the host involved in this alarm. Then, log in to the host as user **omm** and go to [Step 4](#).

**Check the validity period of the certificate files in the system.**

**Step 2** Check whether the remaining validity period of the CA certificate is smaller than the alarm threshold.

Run the **bash \${CONTROLLER\_HOME}/security/cert/conf/querycertvalidity.sh** command to check the effective time and due time of the CA root certificate.

- If yes, go to [Step 5](#).
- If no, go to [Step 7](#).

**Step 3** Check whether the remaining validity period of the HA root certificate is smaller than the alarm threshold.

Run the **openssl x509 -noout -text -in \${CONTROLLER\_HOME}/security/certHA/root-ca.crt** command to check the effective time and due time of the HA root certificate.

- If yes, go to [Step 6](#).
- If no, go to [Step 7](#).

**Step 4** Check whether the remaining validity period of the HA user certificate is smaller than the alarm threshold.

Run the **openssl x509 -noout -text -in \${CONTROLLER\_HOME}/security/certHA/server.crt** command to check the effective time and due time of the HA user certificate.

- If yes, go to [Step 6](#).
- If no, go to [Step 7](#).

The following is an example of the effective time and due time of a CA or HA certificate:

```
Certificate:
Data:
 Version: 3 (0x2)
 Serial Number:
 97:d5:0e:84:af:ec:34:d8
 Signature Algorithm: sha256WithRSAEncryption
 Issuer: C=CN, ST=xxx, L=yyy, O=zzz, OU=IT, CN=HADOOP.COM
 Validity
 Not Before: Dec 13 06:38:26 2016 GMT // Effective time
 Not After : Dec 11 06:38:26 2026 GMT // Due time
```

### Import certificate files.

**Step 5** Import a new CA certificate file.

Apply for or generate a new CA certificate file and import it to the system. For details, see [Replacing the CA Certificate](#). Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to [Step 7](#).
- If no, no further action is required.

**Step 6** Import a new HA certificate file.


Apply for or generate a new HA certificate file and import it to the system. For details, see [Replacing the HA Certificate](#). Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to [Step 7](#).
- If no, no further action is required.

### Collect the fault information.

**Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 8** In the **Services** area, select **Controller**, **OmmServer**, **OmmCore**, and **Tomcat**, and click **OK**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.36 ALM-12057 Metadata Not Configured with the Task to Periodically Back Up Data to a Third-Party Server

### Description

After the system is installed, it checks whether the task for periodically backing up metadata to the third-party server, and then performs the check hourly. If the task for periodically backing up metadata to a third-party server is not configured, a critical alarm is generated.

This alarm is cleared when a user creates such a backup task.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12057	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System


If the metadata is not backed up to a third-party server, when the active/standby management nodes in the cluster are faulty and the local backup data is lost, if you want to use the backup package to restore the cluster metadata, no backup package is available.

### Possible Causes

Metadata is not configured with the task to periodically back up data to a third-party server.

### Procedure

**Step 1** On the FusionInsight Manager portal choose **O&M > Alarm > Alarms**.

**Step 2** In the alarm list, click  in the row where the alarm is located and identify the data module from which the alarm is generated based on **Additional Information**.

**Step 3** Choose **O&M > Backup and Restoration > Backup Management > Create**.

**Step 4** Configure a backup task. The backup data to be configured is consistent with the data in Additional Information of the alarm.

Back up data to a third-party server, for example, remote HDFS (RemoteHDFS), NAS (NFS/CIFS), Object Storage Service (OBS), and SFTP server (SFTP). For details, see [Backing Up Data](#).


**Step 5** After the backup task is created successfully, wait for two minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

#### Collect fault information

**Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 7** In the **Service** area, select **Controller** and click **OK**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.37 ALM-12061 Process Usage Exceeds the Threshold

### Description

The system checks the usage of the omm process every 30 seconds. Users can run the `ps -o nlwp, pid, args, -u omm | awk '{sum+=$1} END {print "", sum}'` command to obtain the number of concurrent processes of user **omm**. Run the `ulimit -u` command to obtain the maximum number of processes that can be simultaneously opened by user **omm**. Divide the number of concurrent processes by the maximum number to obtain the process usage of user **omm**. The process usage has a default threshold. This alarm is generated when the process usage exceeds the threshold.

If **Trigger Count** is **3** and the process usage is less than or equal to the threshold, this alarm is cleared. If **Trigger Count** is greater than **1** and the process usage is less than or equal to 90% of the threshold, this alarm is cleared.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
12061	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Service failure: When the process usage exceeds the threshold, you cannot switch to user **omm**. A new **omm** thread cannot be created. As a result, the job may fail to run.

## Possible Causes

- The alarm threshold is improperly configured.
- The maximum number of processes (including threads) that can be concurrently opened by user **omm** is inappropriate.
- An excessive number of threads are opened at the same time.

## Procedure

**Check whether the alarm threshold or alarm hit number is properly configured.**

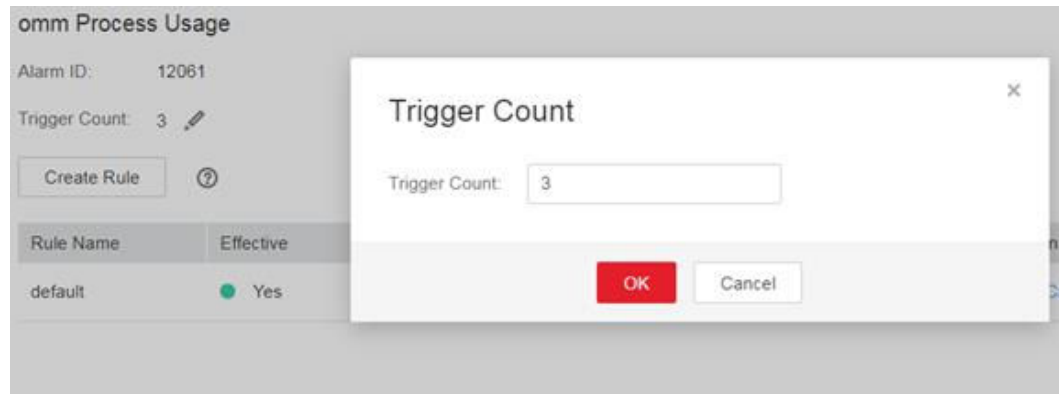
- Step 1** On the FusionInsight Manager, change the alarm threshold and **Trigger Count** based on the actual CPU usage.

Specifically, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Process > omm Process Usage** to change Trigger Count, as shown in [Figure 7-62](#).

**NOTE**

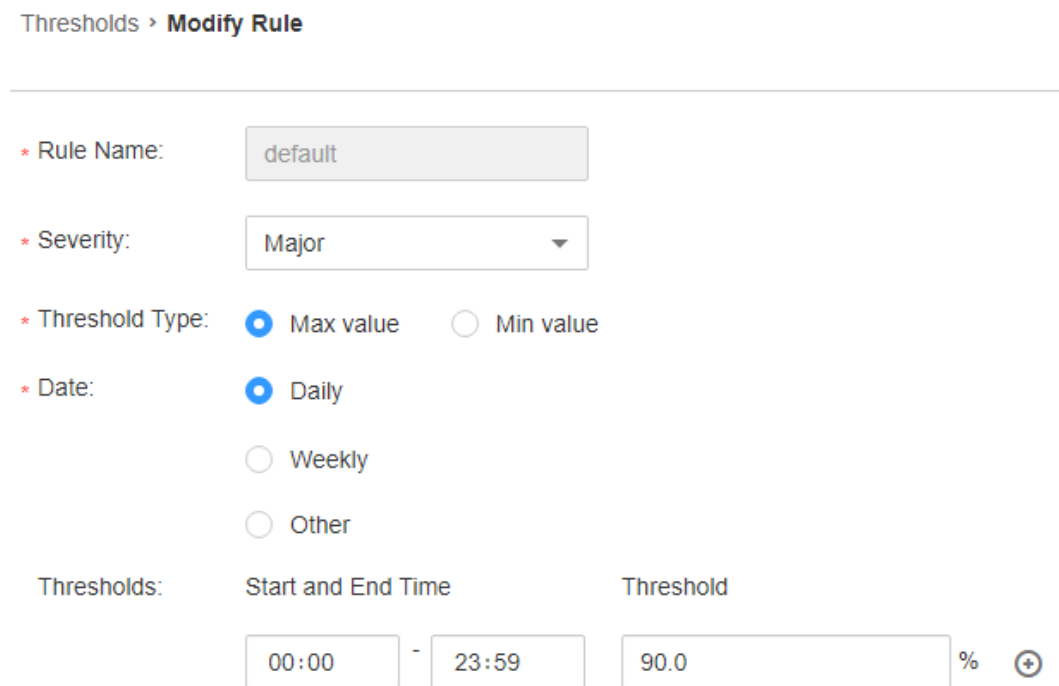
The alarm is generated when the process usage exceeds the threshold for the times specified by **Trigger Count**.

**Figure 7-62** Setting Trigger Count



Set the alarm threshold based on the actual process usage. To check the process usage, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Host > Process > omm Process Usage**, as shown in [Figure 7-63](#).

**Figure 7-63** Setting an alarm threshold



**Step 2** 2 minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 3](#).

**Check whether the maximum number of processes (including threads) opened by user omm is appropriate.**

**Step 3** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host for which the alarm is generated.

**Step 4** Log in to the host where the alarm is generated as user **root**.

**Step 5** Run the **su - omm** command to switch to user **omm**.

**Step 6** Run the **ulimit -u** command to obtain the maximum number of threads that can be concurrently opened by user **omm** and check whether the number is greater than or equal to 60000.

- If it is, go to [Step 8](#).
- If it is not, go to [Step 7](#).

**Step 7** Run the **ulimit -u 60000** command to change the maximum number to 60000. Two minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 12](#).

**Check whether an excessive number of processes are opened at the same time.**

**Step 8** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host for which the alarm is generated.

**Step 9** Log in to the host where the alarm is generated as user **root**.

**Step 10** Run the **ps -o nlwp, pid, lwp, args, -u omm|sort -n** command to check the numbers of threads used by the system.

The result is sorted based on the thread number. Analyze the top 5 thread numbers and check whether the threads are incorrectly used.

- If it is, perform the following operations to stop the parent processes of the top 5 abnormally used processes without affecting services, go to [Step 11](#).
  - a. Run the following command to query the parent process of the corresponding process:  
**ps -ef | grep "Process ID"**  
The third column in the command output is the parent process ID.
  - b. Run the following command to stop the parent process:  
**kill -9 parent process ID**
- If it is not, run the **ulimit -u** command to change the maximum number to be greater than 60000.

**Step 11** Five minutes later, check whether the alarm is cleared.


- If it is, no further action is required.
- If it is not, go to [Step 12](#).

**Collect fault information.**

**Step 12** On the FusionInsight Manager home page of the active clusters, choose **O&M > Log > Download**.

**Step 13** Select **OmmServer** and **NodeAgent** from the **Service** and click **OK**.



**Step 14** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 15** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.38 ALM-12062 OMS Parameter Configurations Mismatch with the Cluster Scale

### Description

The system checks whether the OMS parameter configurations match with the cluster scale at each top hour. If the OMS parameter configurations do not meet the cluster scale requirements, the system generates this alarm. This alarm is automatically cleared when the OMS parameter configurations are modified.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12062	Major	Yes

### Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

If the parameters configured for the current cluster are smaller than the configuration standard required by the cluster scale, problems such as job running delay and slow service page response may occur. In severe cases, the Agent or OMS process on the cluster node is abnormal. As a result, component jobs fail to be submitted and OMS data fails to be synchronized.

## Possible Causes


The OMS parameter configurations mismatch with the cluster scale.

## Procedure

**Check whether the OMS parameter configurations match with the cluster scale.**

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host for which the alarm is generated.
- Step 2** Log in to the host where the alarm is generated as user **root**.
- Step 3** Run the **su - omm** command to switch to user **omm**.
- Step 4** Run the **vi \$BIGDATA\_LOG\_HOME/controller/scriptlog/modify\_manager\_param.log** command to open the log file and search for the log file containing the following information: Current oms configurations cannot support *xx* nodes. In the information, *xx* indicates the number of nodes in the cluster.
- Step 5** Optimize the current cluster configuration by following the instructions in [Optimizing Manager Configurations Based on the Number of Cluster Nodes](#).
- Step 6** One hour later, check whether the alarm is cleared.
  - If it is, no further action is required.
  - If it is not, go to [Step 7](#).

**Collect fault information.**

- Step 7** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 8** Select **Controller** from the **Service** and click **OK**.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

### Optimizing Manager Configurations Based on the Number of Cluster Nodes

**Step 1** Log in to the active Manager node as user **omm**.

**Step 2** Run the following command to switch the directory:

```
cd ${BIGDATA_HOME}/om-server/om/sbin
```

**Step 3** Run the following command to view the current Manager configurations.

```
sh oms_config_info.sh -q
```

**Step 4** Run the following command to specify the number of nodes in the current cluster.

Command format: **sh oms\_config\_info.sh -s *number of nodes***

Example:

```
sh oms_config_info.sh -s 1000
```

Enter **y** as prompted.

The following configurations will be modified:

Module	Parameter	Current	Target
Controller	controller.Xmx	4096m	=> 16384m
Controller	controller.Xms	1024m	=> 8192m
Controller	controller.node.heartbeat.error.threshold	30000	=> 60000
Pms	pms.mem	8192m	=> 10240m

Do you really want to do this operation? (y/n):

The configurations are updated successfully if the following information is displayed:

```
...
Operation has been completed. Now restarting OMS server. [done]
Restarted oms server successfully.
```

#### NOTE

- OMS is automatically restarted during the configuration update process.
- Clusters with similar quantities of nodes have same Manager configurations. For example, when the number of nodes is changed from 100 to 101, no configuration item needs to be updated.

----End

## 7.12.39 ALM-12063 Unavailable Disk

### Description

The system checks whether the data disk of the current host is available at the top of each hour. The system creates files, writes files, and deletes files in the mount directory of the disk. If the operations fail, the alarm is generated. If the operations succeed, the disk is available, and the alarm is cleared.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
12063	Major	Yes

## Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DiskName	Specifies the disk for which the alarm is generated.

## Impact on the System

Service failure: If you need to modify or use data on a disk that is unwritable or unreadable, the job may fail.

## Possible Causes

- The permission of the disk mount directory is abnormal.
- There are disk bad sectors.

## Procedure

**Check whether the permission of the disk mount directory is normal.**

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host and **DiskName** for the disk for which the alarm is generated.
- Step 2** Log in to the host where the alarm is generated as user **root**.
- Step 3** Run the **df -h |grep DiskName** command to obtain the mount point and check whether the permission of the mount directory is unwritable or unreadable.
- If it is, go to [Step 4](#).
  - If it is not, go to [Step 8](#).

 **NOTE**

If the permission of the mount directory is 000 or the owner is **root**, the mount directory is unreadable and unwritable.

**Step 4** Modify the directory permission.

**Step 5** One hour later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

**Step 6** Contact hardware engineers to rectify the disk.


**Step 7** One hour later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 8**.

**Collect fault information.**

**Step 8** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 9** Select **NodeAgent** from the **Service** and click **OK**.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.40 ALM-12064 Host Random Port Range Conflicts with Cluster Used Port

### Alarm Description

The system checks whether the random port range of the host conflicts with the range of ports used by the Cluster system every hour. The alarm is generated if they conflict. The alarm is automatically cleared when the random port range of the host is changed to the normal range.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12064	Major	Yes

## Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

- If ports such as okerberos and oldap are occupied, authentication fails and jobs may fail.
- If the controller and pms ports are occupied, the process is faulty, which may affect the elastic scaling performance.
- If the Tomcat port is occupied, the login and query functions of FusionInsight Manager are affected.

## Possible Causes

The random port range configuration is modified.

## Procedure

### Check the random port range of the system.

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host for which the alarm is generated.
- Step 2** Log in to the host where the alarm is generated as user **root**.
- Step 3** Run the **cat /proc/sys/net/ipv4/ip\_local\_port\_range** command to obtain the random port range of the host and check whether the minimum value is smaller than 32768.
  - If it is, go to [Step 4](#).
  - If it is not, goto [Step 7](#).
- Step 4** Run the **vim /etc/sysctl.conf** command to change the value of **net.ipv4.ip\_local\_port\_range** to **32768 61000**. If this parameter does not exist, add the following configuration: **net.ipv4.ip\_local\_port\_range = 32768 61000**.
- Step 5** Run the **sysctl -p /etc/sysctl.conf** command for the modification to take effect.


**Step 6** One hour later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 7](#).

**Collect fault information.**

**Step 7** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 8** Select **NodeAgent** for **Service** and click **OK**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.41 ALM-12066 Trust Relationships Between Nodes Become Invalid

### Alarm Description

The system checks whether the trust relationship between the active OMS node and other Agent nodes is normal every hour. The alarm is generated if the mutual trust fails. This alarm is automatically cleared after the fault is rectified.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12066	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System


Some operations (such as restart, configuration synchronization, and instance status query) that need to connect to the node may fail. If the trust relationships between nodes are invalid, services may be affected.

## Possible Causes

- The `/etc/ssh/sshd_config` configuration file is damaged.
- The password of user `omm` has expired.

## Handling Procedure

**Check the status of the `/etc/ssh/sshd_config` configuration file.**

**Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm and click  to view the host list in the alarm details.

**Step 2** Log in to the active OMS node as user `omm`.

**Step 3** Run the `ssh` command, for example, `ssh host2`, on each node in the alarm details to check whether the connection fails. (`host2` is a node other than the OMS node in the alarm details.)

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** Open the `/etc/ssh/sshd_config` configuration file on `host2` and check whether `AllowUsers` or `DenyUsers` is configured for other nodes.

- If yes, go to [Step 5](#).
- If no, contact OS experts.

**Step 5** Modify the whitelist or blacklist to ensure that user `omm` is in the whitelist or not in the blacklist. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Check the status of the password of user `omm`.**

**Step 6** Check the interaction information of the `ssh` command.

- If the password of user `omm` is required, go to [Step 7](#).
- If message "Enter passphrase for key '/home/omm/.ssh/id\_rsa':" is displayed, go to [Step 9](#).



- If the connection is set up successfully, run the following command to restart **ssh-agent** and then go to [Step 3](#):

```
ps -ef|grep ssh-agent |grep -v grep |awk '{print $2}' |xargs kill -9
```

**Step 7** Check the trust list (`/home/omm/.ssh/authorized_keys`) of user **omm** on the OMS node and `host2` node. Check whether the trust list contains the public key file (`/home/omm/.ssh/id_rsa.pub`) of user **omm** on the peer host.

- If yes, contact OS experts.
- If no, add the public key of user **omm** of the peer host to the trust list of the local host.


**Step 8** Add the public key of user **omm** of the peer host to the trust list of the local host. Run the **ssh** command, for example, **ssh host2**, on each node in the alarm details to check whether the connection fails. (*host2* is a node other than the OMS node in the alarm details.)

- If yes, go to [Step 9](#).
- If no, check whether the alarm is cleared. If the alarm is cleared, no further action is required; otherwise, go to [Step 9](#).

#### Collect the fault information.

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Select **Controller** for **Service** and click **OK**.

**Step 11** Click  in the upper right corner to set the log collection time range. Generally, the time range is 10 minutes before and after the alarm generation time. Click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

Perform the following steps to handle abnormal trust relationships between nodes:

---

### NOTICE

- Perform this operation as user **omm**.
- If the network between nodes is disconnected, rectify the network fault first. Check whether the two nodes are connected to the same security group and whether **hosts.deny** and **hosts.allow** are set.

1. Run the **ssh-add -l** command on both nodes to check whether any identities exist.

```

omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ll .ssh/
total 32
-rw-----, 1 omm wheel 0 Dec 29 14:17 agent.pid
-rw-----, 1 omm wheel 12901 Mar 9 14:48 authorized_keys
-rw-----, 1 omm wheel 54 Sep 24 11:42 config
-rw-----, 1 omm wheel 1766 Sep 24 11:43 id_rsa
-rw-----, 1 omm wheel 402 Sep 24 11:42 id_rsa.pub
-rw-----, 1 omm wheel 88 Jun 8 2020 id_rsa.sha256
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ssh-add -l
The agent has no identities.
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ vim /var/log/Bigdata/nodeagent/
agentlog/ alarmlog/ monitorlog/ scriptlog/
omm@node-group-2eU40 ~]$ vim /var/log/Bigdata/nodeagent/scriptlog/
agent_alarm_py.log install_log
agent_alarm_py.log.1 installntp.log

```

- If yes, go to 4.
  - If no, go to 2.
2. If no identities are displayed, run the **ps -ef|grep ssh-agent** command to find the **ssh-agent** process, stop the process, and wait for the process to automatically restart.

```

omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ssh-add -l
The agent has no identities.
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ps -ef|grep ssh-agent
omm 18729 1 0 14:53 ? 00:00:00 ssh-agent -a /home/omm/.ssh/agent.pid
omm 25098 1 0 14:54 ? 00:00:00 bash /opt/Bigdata/om-agent/nodeagent/bin/ssh-agent-monitor-startup.sh
omm 25286 25098 0 14:54 ? 00:00:00 bash /opt/Bigdata/om-agent/nodeagent/bin/ssh-agent-monitor.sh
omm 27201 4913 0 14:54 pts/0 00:00:00 grep --color=auto ssh-agent
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ssh-add -l

```

3. Run the **ssh-add -l** command to check whether the identities have been added. If yes, manually run the **ssh** command to check whether the trust relationship is normal.

```

omm 22276 4913 0 14:53 pts/0 00:00:00 grep --color=auto ssh-agent
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ssh-add -l
The agent has no identities.
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ps -ef|grep ssh-agent
omm 18729 1 0 14:53 ? 00:00:00 ssh-agent -a /home/omm/.ssh/agent.pid
omm 25098 1 0 14:54 ? 00:00:00 bash /opt/Bigdata/om-agent/nodeagent/bin/ssh-agent-monitor-startup.sh
omm 25286 25098 0 14:54 ? 00:00:00 bash /opt/Bigdata/om-agent/nodeagent/bin/ssh-agent-monitor.sh
omm 27201 4913 0 14:54 pts/0 00:00:00 grep --color=auto ssh-agent
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ssh-add -l
2048 SHA256:uchnRUBhh1HxptF0Z1bS0zym1UKXm1afYvn0IMpiZjg /home/omm/.ssh/id_rsa (RSA)
omm@node-group-2eU40 ~]$
omm@node-group-2eU40 ~]$ ssh 10.33.109.226
Warning: Permanently added '10.33.109.226' (ECDSA) to the list of known hosts.

```

4. If identities exist, check whether the **/home/omm/.ssh/authorized\_keys** file contains the information in the **/home/omm/.ssh/id\_rsa.pub** file of the peer node. If it does not, manually add the information.
5. Check whether the permissions on the files in the **/home/omm/.ssh** directory are modified.
6. Check the **/var/log/Bigdata/nodeagent/scriptlog/ssh-agent-monitor.log** file.
7. If the **/home** directory of user **omm** is deleted, contact MRS support personnel for assistance.

## 7.12.42 ALM-12067 Tomcat Resource Is Abnormal

### Alarm Description

HA checks the Tomcat resources of Manager every 85 seconds. This alarm is generated when HA detects that the Tomcat resources are abnormal for two consecutive times.

This alarm is cleared when HA detects that the Tomcat resources become normal.

**Resource Type** of Tomcat is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new Tomcat resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12067	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System


An active/standby Manager switchover may occur. The Manager and component web UIs are unavailable. The cluster management function cannot be provided for upper-layer web applications, and users may fail to log in to the Manager and component web UIs.

### Possible Causes


- The Tomcat directory permission is abnormal, and the Tomcat process is abnormal.

## Handling Procedure

**Check whether the permission on the Tomcat directory is normal.**

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click  to view the IP address of the host for which the alarm is generated.
- Step 2** Log in to the alarm host as user **root**.
- Step 3** Run the **su - omm** command to switch to user **omm**.
- Step 4** Run the **vi \$BIGDATA\_LOG\_HOME/omm/oms/ha/scriptlog/tomcat.log** command to check whether the Tomcat resource log contains keyword **Cannot find XXX** and rectify the file permission based on the keyword.
- Step 5** After 5 minutes, check whether the alarm is automatically cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** In the **Services** area, select **OmmServer** and **Tomcat**, and click **OK**.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.43 ALM-12068 ACS Resource Exception

### Alarm Description

HA checks the ACS resources of Manager every 80 seconds. This alarm is generated when HA detects that the ACS resources are abnormal for two consecutive times.

This alarm is cleared when HA detects that the ACS resources are normal.

**Resource Type** of ACS is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new ACS resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12068	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System


An active/standby Manager switchover may occur. The security authentication and user management functions cannot be provided for ACS upper-layer applications. As a result, you may fail to log in to Manager and component web UIs.

## Possible Causes

The ACS process is abnormal.

## Handling Procedure

**Check whether the ACS process is normal.**

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click  to view the name of the host for which the alarm is generated.
- Step 2** Log in to the alarm host as user **root**.
- Step 3** Run the **su - omm** command and then **sh \${BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** to check whether the status of the ACS resources managed by the HA is normal. In the single-node system, the ACS resource is in the normal state. In the dual-node system, the ACS resource is in the normal state on the active node and in the stopped state on the standby node.
- If yes, go to [Step 6](#).
  - If no, go to [Step 4](#).

**Step 4** Run the `vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/acs.log` command to check whether the ACS resource log of HA contains the keyword **ERROR**. If yes, analyze the logs to locate the resource exception cause and fix the exception.


**Step 5** After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect the fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** In the **Services** area, select **Controller** and **OmmServer**, and click **OK**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.44 ALM-12069 AOS Resource Exception

### Alarm Description

HA checks the AOS resources of Manager every 81 seconds. This alarm is generated when HA detects that the AOS resources are abnormal for two consecutive times.

This alarm is cleared when HA detects that the AOS resources become normal.

**Resource Type** of AOS is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new AOS resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12069	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System


An active/standby Manager switchover may occur. Tenant and role management cannot be provided for AOS upper-layer applications. As a result, you may fail to log in to Manager and component web UIs.

## Possible Causes


The AOS process is abnormal.

## Handling Procedure

**Check whether the AOS process is normal.**

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click  to view the name of the host for which the alarm is generated.
- Step 2** Log in to the alarm host as user **root**.
- Step 3** Run the **su - omm** command and then **sh \${BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** to check whether the status of the AOS resources managed by the HA is normal. In the single-node system, the AOS resource is in the normal state. In the dual-node system, the AOS resource is in the normal state on the active node and in the stopped state on the standby node.
  - If yes, go to [Step 6](#).
  - If no, go to [Step 4](#).
- Step 4** Run the **vi \$BIGDATA\_LOG\_HOME/omm/oms/ha/scriptlog/aos.log** command to check whether the AOS resource log of HA contains the keyword **ERROR**. If yes, analyze the logs to locate the resource exception cause and fix the exception.
- Step 5** After 5 minutes, check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 6](#).

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** In the **Services** area, select **Controller** and **OmmServer**, and click **OK**.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.45 ALM-12070 Controller Resource Is Abnormal

### Alarm Description

HA checks the controller resources of Manager every 80 seconds. This alarm is generated when HA detects that the controller resources are abnormal for 2 consecutive times.

This alarm is cleared when the Controller resource is normal.

**Resource Type** of Controller is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new Controller resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12070	Major	Yes

### Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.



Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

- The alarm persists for a long time, causing frequent active/standby switchovers of FusionInsight Manager. As a result, users cannot log in to FusionInsight Manager and perform O&M operations.
- The Controller process repeatedly restarts, which may cause the native UI of the service login failure.

## Possible Causes

The Controller process is abnormal.

## Procedure


**Check whether the controller process is normal.**

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.
- Step 2** Log in to the host for which the alarm is generated as user **root**.
- Step 3** Run the **su - omm** command to switch to user **omm**. Run the **sh \$ {BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** command to check whether the status of the Controller resources managed by the HA is normal. In the single-node system, the Controller resource is in the normal state. In the dual-node system, the Controller resource is in the normal state on the active node and in the stopped state on the standby node.
- If it is, go to **Step 6**.
  - If it is not, go to **Step 4**.
- Step 4** Run the **vi \$BIGDATA\_LOG\_HOME/omm/oms/ha/scriptlog/controller.log** command to view the Controller resource logs, and run the **vi \$BIGDATA\_LOG\_HOME/controller/controller.log** command to view the Controller running logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.
- Step 5** Five minutes later, check whether this alarm is cleared.
- If it is, no further action is required.
  - If it is not, go to **Step 6**.

**Collect fault information.**

- Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 7** Select **Controller** and **OmmServe** for **Service** and click **OK**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour before and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.46 ALM-12071 Httpd Resource Is Abnormal

### Description

HA checks the httpd resources of Manager every 120 seconds. This alarm is generated when HA detects that the httpd resources are abnormal for 10 consecutive times.

This alarm is cleared when the httpd resource is normal.

**Resource Type** of httpd is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new httpd resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12071	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

- The alarm persists for a long time, causing frequent active/standby switchovers of FusionInsight Manager. As a result, users cannot log in to FusionInsight Manager and perform O&M operations.
- The httpd process is repeatedly restarts, which may lead to the failure to visit the native service UI.

## Possible Causes


The httpd process is abnormal.

## Procedure

### Check whether the httpd process is abnormal.

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.
- Step 2** Log in to the host for which the alarm is generated as user **root**.
- Step 3** Run the **su - omm** command to switch to user **omm**.
- Step 4** Run the **sh \${BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** command to check whether the status of the httpd resources managed by the HA is normal. In the single-node system, the httpd resource is in the normal state. In the dual-node system, the httpd resource is in the normal state on the active node and in the stopped state on the standby node.
- If it is, go to [Step 7](#).
  - If it is not, go to [Step 5](#).
- Step 5** Run the **vi \$BIGDATA\_LOG\_HOME/omm/oms/ha/scriptlog/httpd.log** command to view the httpd resource logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.
- Step 6** Five minutes later, check whether this alarm is cleared.
- If it is, no further action is required.
  - If it is not, go to [Step 7](#).

### Collect fault information.

- Step 7** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 8** Select **Controller** and **OmmServer** for **Service** and click **OK**.
- Step 9** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.47 ALM-12072 FloatIP Resource Is Abnormal

### Description

HA checks the floatip resources of Manager every 9 seconds. This alarm is generated when HA detects that the floatip resources are abnormal for 3 consecutive times.

This alarm is cleared when the FloatIP resource is normal.

**Resource Type** of FloatIP is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new FloatIP resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12072	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

- The alarm persists for a long time, causing frequent active/standby switchovers of FusionInsight Manager. As a result, users cannot log in to FusionInsight Manager and perform O&M operations.
- The FloatIP process is repeatedly restarts, which may lead to the failure to visit the native service UI.

## Possible Causes

- The floating IP address is abnormal.

## Procedure

### Check the floating IP address status of the active management node.

**Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the address of the host for which the alarm is generated and the resource name.

**Step 2** Log in to the active management node as user **root**.

**Step 3** Run the following command, go to the `/${BIGDATA_HOME}/om-server/om/sbin/` directory.

```
su - omm
```

```
cd ${BIGDATA_HOME}/om-server/om/sbin/
```

**Step 4** Run the `sh status-oms.sh` command, and execute the `status-oms.sh` script to check whether the floating IP address of the active FusionInsight Manager is normal. View the command output, locate the row where **ResName** is **floatip**, and check whether the following information is displayed.

For example:

```
10-10-10-160 floatip Normal Normal Single_active
```

- If it is, go to [Step 8](#).
- If it is not, go to [Step 5](#).

**Step 5** Run the `ifconfig` command to check whether the NIC with the floating IP address exists.

- If it does, go to [Step 8](#).
- If it does not, go to [Step 6](#).


**Step 6** Run the `ifconfig NIC name Floating IPaddress netmask Subnet mask` command to reconfigure the NIC with the floating IP address. (For example, `ifconfig eth0 10.10.10.102 netmask 255.255.255.0`).

**Step 7** Five minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 8](#).

### Collect fault information.

**Step 8** On FusionInsight Manager, choose **O&M > Log > Download**.

- Step 9** Select **Controller** and **OmmServer** for **Service** and click **OK**.
- Step 10** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.
- Step 11** Contact the O&M personnel and send the collected log information.
- End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.48 ALM-12073 CEP Resource Is Abnormal

### Description

HA checks the cep resources of Manager every 60 seconds. This alarm is generated when HA detects that the cep resources are abnormal for 2 consecutive times.

This alarm is cleared when the CEP resource is normal.

**Resource Type** of CEP is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new CEP resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12073	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

- The alarm persists for a long time, causing frequent active/standby switchovers of FusionInsight Manager. As a result, users cannot log in to FusionInsight Manager and perform O&M operations.
- The CEP process repeatedly restarts. As a result, monitoring data collection fails during the alarm reporting period. In severe cases, monitoring data during the alarm reporting period may be lost.

## Possible Causes

The CEP process is abnormal.


## Procedure

### Check whether the CEP process is abnormal.

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.
- Step 2** Log in to the host for which the alarm is generated as user **root**.
- Step 3** Run the **su -omm** command and then the **sh \${BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** command to check whether the status of the CEP resources managed by the HA is normal. In the single-node system, the CEP resource is in the normal state. In the dual-node system, the CEP resource is in the normal state on the active node and in the stopped state on the standby node.
- If it is, go to [Step 6](#).
  - If it is not, go to [Step 4](#).
- Step 4** Run the **vi \$BIGDATA\_LOG\_HOME/omm/oms/cep/cep.log** and **vi \$BIGDATA\_LOG\_HOME/omm/oms/cep/scriptlog/cep\_ha.log** commands to view the CEP resource logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.
- Step 5** Five minutes later, check whether this alarm is cleared.
- If it is, no further action is required.
  - If it is not, go to [Step 6](#).

### Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 7** Select **Controller** and **OmmServer** for **Service** and click **OK**.

**Step 8** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.49 ALM-12074 FMS Resource Is Abnormal

### Description

HA checks the fms resources of Manager every 60 seconds. This alarm is generated when HA detects that the fms resources are abnormal for 2 consecutive times.

This alarm is cleared when the FMS resource is normal.

**Resource Type** of FMS is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new FMS resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12074	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.



Name	Meaning
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

- The alarm persists for a long time, causing frequent active/standby switchovers of FusionInsight Manager. As a result, users cannot log in to FusionInsight Manager and perform O&M operations.
- The FMS process repeatedly restarts. As a result, the alarm data reported during the alarm reporting period is abnormal. In severe cases, the alarm data reported during the alarm reporting period may fail to be reported and cleared.

## Possible Causes

The FMS process is abnormal.


## Procedure

### Check whether the FMS process is abnormal.

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.
- Step 2** Log in to the host for which the alarm is generated as user **root**.
- Step 3** Run the **su -omm** command and then the **sh \${BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** command to check whether the status of the FMS resources managed by the HA is normal. In the single-node system, the FMS resource is in the normal state. In the dual-node system, the FMS resource is in the normal state on the active node and in the stopped state on the standby node.
- If it is, go to [Step 6](#).
  - If it is not, go to [Step 4](#).
- Step 4** Run the **vi \$BIGDATA\_LOG\_HOME/omm/oms/fms/fms.log** and **vi \$BIGDATA\_LOG\_HOME/omm/oms/fms/scriptlog/fms\_ha.log** commands to view the FMS resource logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.
- Step 5** 5 minutes later, check whether this alarm is cleared.
- If it is, no further action is required.
  - If it is not, go to [Step 6](#).

### Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M> Log > Download**.
- Step 7** Select **Controller** and **OmmServer** for **Service** and click **OK**.

**Step 8** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.50 ALM-12075 PMS Resource Is Abnormal

### Description

HA checks the pms resources of Manager every 55 seconds. This alarm is generated when HA detects that the pms resources are abnormal for three consecutive times.

This alarm is cleared when the PMS resource is normal.

**Resource Type** of PMS is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new PMS resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12075	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

- The alarm persists for a long time, causing frequent active/standby switchovers of FusionInsight Manager. As a result, users cannot log in to FusionInsight Manager and perform O&M operations.
- The PMS process repeatedly restarts. As a result, monitoring data collection fails during the alarm reporting period. In severe cases, monitoring data may be lost during the alarm reporting period.

## Possible Causes

The PMS process is abnormal.


## Procedure

### Check whether the PMS process is abnormal.

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.
- Step 2** Log in to the host for which the alarm is generated as user **root**.
- Step 3** Run the **su -omm** command and then the **sh \${BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** command to check whether the status of the PMS resources managed by the HA is normal. In the single-node system, the PMS resource is in the normal state. In the dual-node system, the PMS resource is in the normal state on the active node and in the stopped state on the standby node.
- If it is, go to [Step 6](#).
  - If it is not, go to [Step 4](#).
- Step 4** Run the **vi \$BIGDATA\_LOG\_HOME/omm/oms/pms/pms.log** and **vi \$BIGDATA\_LOG\_HOME/omm/oms/pms/scriptlog/pms\_ha.log** commands to view the PMS resource logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.
- Step 5** Five minutes later, check whether this alarm is cleared.
- If it is, no further action is required.
  - If it is not, go to [Step 6](#).

### Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 7** Select **Controller** and **OmmServer** for **Service** and click **OK**.

**Step 8** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.51 ALM-12076 GaussDB Resource Is Abnormal

### Description

HA checks the Manager database every 10 seconds. This alarm is generated when HA detects that the database is abnormal for 3 consecutive times.

This alarm is cleared when the database is normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12076	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

If the database is abnormal, all core services and related service processes of Manager, such as the alarm, monitoring, and query functions, are affected.

## Possible Causes

An exception occurs in the database.

## Procedure

**Check the database status of the active and standby management nodes.**

- Step 1** Log in to the active and standby management nodes respectively as user **root**. Run the **su - ommdba** command to switch to user **ommdba**, and then run the **gs\_ctl query** command to check whether the following information is displayed in the command output.

Command output of the active management node:

```
Ha state:
LOCAL_ROLE: Primary
STATIC_CONNECTIONS : 1
DB_STATE : Normal
DETAIL_INFORMATION : user/password invalid
Senders info:
No information
Receiver info:
No information
```

Command output of the standby management node:

```
Ha state:
LOCAL_ROLE: Standby
STATIC_CONNECTIONS : 1
DB_STATE : Normal
DETAIL_INFORMATION : user/password invalid
Senders info:
No information
Receiver info:
No information
```

- If it is, go to [Step 3](#).
- If it is not, go to [Step 2](#).

- Step 2** Contact the network administrator to check whether the network is faulty.

- If it is, go to [Step 3](#).
- If it is not, go to [Step 5](#).

- Step 3** Five minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

- Step 4** Log in to the active and standby management nodes, run the **su -omm** command to switch to user **omm**, go to the **`\${BIGDATA\_HOME}/om-server/om/sbin/`** directory, and run the **status-oms.sh** script to check whether the floating IP addresses and GaussDB resources of the active and standby FusionInsight Managers are in the status shown in the following figure.


acs	Normal	Normal	Single_active
aos	Normal	Normal	Single_active
cep	Normal	Normal	Single_active
controller	Normal	Normal	Single_active
feed_watchdog	Normal	Normal	Double_active
floatip	Normal	Normal	Single_active
fms	Normal	Normal	Single_active
gaussDB	Active_normal	Normal	Active_standby
heartBeatCheck	Normal	Normal	Single_active
httpd	Normal	Normal	Single_active
iam	Normal	Normal	Single_active
ntp	Active_normal	Normal	Active_standby
okerberos	Normal	Normal	Double_active
oldap	Active_normal	Normal	Active_standby
pms	Normal	Normal	Single_active
tomcat	Normal	Normal	Single_active
acs	Stopped	Normal	Single_active
aos	Stopped	Normal	Single_active
cep	Stopped	Normal	Single_active
controller	Stopped	Normal	Single_active
feed_watchdog	Normal	Normal	Double_active
floatip	Stopped	Normal	Single_active
fms	Stopped	Normal	Single_active
gaussDB	Standby_normal	Normal	Active_standby
heartBeatCheck	Stopped	Normal	Single_active
httpd	Stopped	Normal	Single_active

- If they are, find the alarm in the alarm list and manually clear the alarm.
- If they are not, go to [Step 5](#).

#### Collect fault information.

**Step 5** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 6** Select **OmmServer** for **Service** and click **OK**.

**Step 7** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 8** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.52 ALM-12077 User omm Expired

### Description

The system starts at 00:00 every day to check whether user **omm** has expired every eight hours. This alarm is generated if the user account has expired.

This alarm is cleared when the expiration time of user **omm** is changed and the user account status becomes normal.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
12077	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

User **omm** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services.

## Possible Causes

User **omm** has expired.

## Procedure

**Check whether user omm in the system has expired.**

**Step 1** Log in to the faulty node as user **root**.

Run the **chage -l omm** command to view the information about the password of user **omm**.

**Step 2** View the value of **Account expires** to check whether the user configurations have expired.

### NOTE

If the parameter value is **never**, the user configurations never expire.

- If they do, go to [Step 3](#).
- If they do not, go to [Step 4](#).


**Step 3** Run the **chage -E 'yyyy-MM-dd' omm** command to set the expiration time of user **omm**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 5** Select **NodeAgent** for **Service** and click **OK**.

**Step 6** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.53 ALM-12078 Password of User omm Expired

### Description

The system starts at 00:00 every day to check whether the password of user **omm** has expired every 8 hours. This alarm is generated if the password has expired.

This alarm is cleared when the expiration time of user **omm** password is changed and the user password status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12078	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.



Name	Meaning
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

The password of user **omm** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services. The **crontab** scheduled task cannot be executed, affecting the ClickHouse service.

## Possible Causes

The password of user **omm** has expired.

## Procedure

**Check whether the password of user omm in the system has expired.**

**Step 1** Log in to the faulty node as user **root**.

Run the **chage -l omm** command to view the information about the password of user **omm**.

**Step 2** View the value of **Password expires** to check whether the user configurations have expired.

### NOTE

If the parameter value is **never**, the user configurations never expire.

- If they do, go to [Step 3](#).
- If they do not, go to [Step 4](#).


**Step 3** Run the **chage -M 'days' omm** command to set the validity period of the password for user **omm**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M> Log > Download**.

**Step 5** Select **NodeAgent** for **Service** and click **OK**.

**Step 6** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.54 ALM-12079 User omm Is About to Expire

### Description

The system starts at 00:00 every day to check whether user **omm** is about to expire every 8 hours. This alarm is generated if the user account will expire no less than 15 days later.

This alarm is cleared when the expiration time of user **omm** is changed and the user account status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12079	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

User **omm** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services.

### Possible Causes

The account of user **omm** is about to expire.

## Procedure

### Check whether user omm is about to expire.

**Step 1** Log in to the faulty node as user **root**.

Run the **chage -l omm** command to view the information about the password of user **omm**.

**Step 2** View the value of **Account expires** to check whether the user configurations are about to expire.

#### NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If they are, go to [Step 3](#).
- If they are not, go to [Step 4](#).


**Step 3** Run the **chage -E 'yyyy-MM-dd' omm** command to set the validity period of user **omm**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

### Collect fault information.

**Step 4** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 5** Select **NodeAgent** for **Service** and click **OK**.

**Step 6** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.55 ALM-12080 Password of User omm Is About to Expire

### Description

The system starts at 00:00 every day to check whether the password of user **omm** is about to expire every 8 hours. This alarm is generated if the password will expire no less than 15 days later.

This alarm is cleared when the expiration time of user **omm** password is reset and the user password status becomes normal.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
12080	Minor	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

The password of user **omm** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services. The **crontab** scheduled task cannot be executed, affecting the ClickHouse service.

## Possible Causes

The password of user **omm** is about to expire.

## Procedure

**Check whether the password of user omm in the system is about to expire.**

**Step 1** Log in to the faulty node as user **root**.

Run the **chage -l omm** command to view the information about the password of user **omm**.

**Step 2** View the value of **Password expires** to check whether the user configurations are about to expire.

### NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If they are, go to [Step 3](#).
- If they are not, go to [Step 4](#).


**Step 3** Run the **chage -M 'days' omm** command to set the validity period of the password for user **omm**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M> Log > Download**.

**Step 5** Select **NodeAgent** for **Service** and click **OK**.

**Step 6** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.56 ALM-12081 User ommdba Expired

### Description

The system starts at 00:00 every day to check whether user **ommdba** has expired every 8 hours. This alarm is generated if the user account has expired.

This alarm is cleared when the expiration time of user **ommdba** is reset and the user account status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12081	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

The OMS database cannot be managed and data cannot be accessed.

## Possible Causes

The account of user **ommdba** for the host has expired.

## Procedure

### Check whether user **ommdba** has expired.

**Step 1** Log in to the faulty node as user **root**.

Run the **chage -l ommdba** command to view the information about the password of user **ommdba**.

**Step 2** View the value of **Account expires** to check whether the user configurations have expired.

#### NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password have expired.

- If they do, go to [Step 3](#).
- If they do not, go to [Step 4](#).


**Step 3** Run the **chage -E 'yyyy-MM-dd' omm** command to set the validity period of user **ommdba**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

### Collect fault information.

**Step 4** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 5** Select **NodeAgent** for **Service** and click **OK**.

**Step 6** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.57 ALM-12082 User ommdba Is About to Expire

### Description

The system starts at 00:00 every day to check whether user **ommdba** is about to expire every 8 hours. This alarm is generated if the user account will expire no less than 15 days later.

This alarm is cleared when the expiration time of user **ommdba** is reset and the user account status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12082	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

The OMS database cannot be managed and data cannot be accessed.

## Possible Causes

The account of user **ommdba** for the host is about to expire.

## Procedure

**Check whether user ommdba is about to expire.**

**Step 1** Log in to the faulty node as user **root**.

Run the **chage -l ommdba** command to view the information about user **ommdba**.

**Step 2** View the value of **Account expires** to check whether the user configurations are about to expire.

### NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If they are, go to [Step 3](#).
- If they are not, go to [Step 4](#).


**Step 3** Run the **chage -E 'yyyy-MM-dd' ommdba** command to set the validity period of user **ommdba**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 5** Select **NodeAgent** for **Service** and click **OK**.

**Step 6** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None



## 7.12.58 ALM-12083 Password of User ommdba Is About to Expire

### Description

The system starts at 00:00 every day to check whether the password of user **ommdba** is about to expire every 8 hours. This alarm is generated if the password is about to expire no less than 15 days later.

This alarm is cleared when the expiration time of user **ommdba** password is reset and the user password status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12083	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

The OMS database cannot be managed and data cannot be accessed.

### Possible Causes

The password of user **ommdba** is about to expire.

### Procedure

**Check whether the password of user ommdba in the system is about to expire.**

**Step 1** Log in to the faulty node as user **root**.

Run the **chage -l ommdba** command to view the information about the password of user **ommdba**.

**Step 2** View the value of **Password expires** to check whether the user configurations are about to expire.

 **NOTE**

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If they are, go to **Step 3**.
- If they are not, go to **Step 4**.


**Step 3** Run the **chage -M 'days' ommdba** command to set the validity period of the password for user **ommdba**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 5** Select **NodeAgent** for **Service** and click **OK**.

**Step 6** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.59 ALM-12084 Password of User ommdba Expired

### Description

The system starts at 00:00 every day to check whether the password of user **ommdba** has expired every 8 hours. This alarm is generated if the password has expired.

This alarm is cleared when the expiration time of user **ommdba** password is reset and the user password status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12084	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

The password of user **ommdba** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services.

## Possible Causes

The password of user **ommdba** for the host has expired.

## Procedure

**Check whether the password of user ommdba in the system has expired.**

**Step 1** Log in to the faulty node as user **root**.

Run the **chage -l ommdba** command to view the information about the password of user **ommdba**.

**Step 2** View the value of **Password expires** to check whether the user configurations have expired.

 **NOTE**

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password have expired.

- If they do, go to [Step 3](#).
- If they do not, go to [Step 4](#).


**Step 3** Run the **chage -M 'days' ommdba** command to set the validity period of the password for user **ommdba**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 5** Select **NodeAgent** for **Service** and click **OK**.

**Step 6** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 7** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.60 ALM-12085 Service Audit Log Dump Failure

### Description

The system dumps service audit logs at 03:00 every day and stores them on the OMS node. This alarm is generated when the dump fails. This alarm is cleared when the next dump succeeds.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12085	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

If the audit logs of a component fail to be dumped, the audit logs cannot be retrieved if they are aged locally. This affects service analysis and troubleshooting of the component.

## Possible Causes

- The service audit logs are oversized.
- The OMS backup storage space is insufficient.
- The storage space of a host where the service is located is insufficient.

## Procedure

### Check whether the service audit logs are oversized.

**Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host and additional information for which the alarm is generated.

**Step 2** Log in to the host where the alarm is generated as user **root**.

**Step 3** Run the **vi \${BIGDATA\_LOG\_HOME}/controller/scriptlog/getLogs.log** command to check whether the keyword "LOG SIZE is more than 5000MB" can be searched.

- If it can, go to [Step 4](#).
- If it cannot, go to [Step 5](#).

**Step 4** Check whether the oversized service audit logs are caused by exceptions.

### The OMS backup storage space is insufficient.

**Step 5** Run the **vi \${BIGDATA\_LOG\_HOME}/controller/scriptlog/getLogs.log** command to check whether the keyword "Collect log failed, too many logs on" can be searched.

- If it can, obtain the host IP address following the keyword "Collect log failed, too many logs on", and go to [Step 6](#).
- If it cannot, go to [Step 11](#).

**Step 6** Log in to the host with the IP address obtained in [Step 5](#) as user **root**.

**Step 7** Run the **vi \${BIGDATA\_LOG\_HOME}/nodeagent/scriptlog/collectLog.log** command to check whether the keyword "log size exceeds" can be searched.

- If it can, go to [Step 9](#).
- If it cannot, go to [Step 8](#).

**Step 8** Check whether the alarm additional information contains the keyword "no enough space".

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

**Step 9** Perform the following operations to expand the disk capacity (only for MRS 3.1.2 and earlier versions) or reduce the maximum number of audit log backups:

- Expand the capacity of the OMS node.

- Run the following command to edit the file and decrease the value of **MAX\_NUM\_BK\_AUDITLOG**.

```
vi ${CONTROLLER_HOME}/etc/om/componentsauditlog.properties
```

**Step 10** In the next execution period, 03:00, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 11](#).

**Check whether the space of the host where the service is located is insufficient.**

**Step 11** Run the `vi ${BIGDATA_LOG_HOME}/controller/scriptlog/getLogs.log` command to check whether the keyword "Collect log failed, no enough space on *hostip*" can be searched.

- If it can, obtain the IP address of the abnormal host and go to [Step 12](#).
- If it cannot, go to [Step 15](#).

**Step 12** Log in to the host with the IP address obtained as user **root**, and run the `df "$BIGDATA_HOME/tmp" -lP | tail -1 | awk '{print ($4/1024)}'` command to obtain the remaining space of the host log directory. Check whether the value is less than 1000 MB.

- If it is, go to [Step 13](#).
- If it is not, go to [Step 15](#).

**Step 13** Expand the capacity of the node


**Step 14** In the next execution period, 03:00, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 15](#).

**Collect fault information.**

**Step 15** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 16** Select **Controller** for **Service** and click **OK**.

**Step 17** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 18** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.61 ALM-12087 System Is in the Upgrade Observation Period

### Description

The system checks whether it is in the upgrade observation period at 00:00 every day and checks whether the duration that it has been in the upgrade observation state exceeds the preset upgrade observation period, 10 days by default. This alarm is generated when the system is in the upgrade observation period and the duration that the system has been in the upgrade observation state exceeds the preset period (10 days by default). This alarm is automatically cleared if the system exits the upgrade observation period after the user performs a rollback or submission.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12087	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Upgrade Observation Period (Days)	Specifies the days that the system is in the upgrade observation period.

### Impact on the System

During the upgrade observation period, do not add or delete users, instances, roles, services, hosts, or resource pools that affect the management topology.

### Possible Causes

The upgrade task is not submitted a specified period of time (10 days by default) after the system upgrade.

## Procedure

**Check whether the system is in the upgrade observation period.**

**Step 1** Log in to the active management node as user **root**.

**Step 2** Run the following commands to switch to user **omm** and log in to the **omm** database:

```
su - omm
```

```
gsqll -U omm -W omm database password -p 20015
```

**Step 3** Run the **select \* from OM\_CLUSTERS** command to view cluster information.

**Step 4** Check whether the value of **upgradObservationPeriod isON** is **true**, as shown in [Figure 7-64](#).

- If it is, the system is in the upgrade observation period. Use the UpdateTool to submit the upgrade task. For details, see the upgrade guide of the corresponding version.
- If it is not, go to [Step 6](#).

**Figure 7-64** Cluster information

```
CLUSTER_ID | CLUSTER_NAME | CLUSTER_DESCRIPTION | STACK_NAME | STACK_TIME | PRESTACK_NAME | PRESTACK_TIME | STACK_MODEL | CURRENT_PATCH_VERSION | IS_DETACHED | UPDATE_MODE | OBSERVATION_PERIOD | EXTERNAL_PARAM |
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
cluster_1 | Test_1 | | DEFAULT_STACK | 1552290738686 | | Sec | | 0 | | {'upgradObservationPeriod': {'isOn': true, 'proje |
': '199318993146781010', 'type': 'UPGRADE'}, 'updateEndTime': '1552301484884', 'patchObservationPeriod': {'isOn': false, 'updateEndTime': '03' | }
```


**Step 5** In the early morning of the next day, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 6](#).

**Collect fault information.**

**Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 7** Select **Controller** from the **Service** and click **OK**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None



## 7.12.62 ALM-12089 Inter-Node Network Is Abnormal

### Description

The alarm module checks the network health status of nodes in the cluster every 10 seconds. This alarm is generated when the network between two nodes is unreachable or the network status is unstable.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12089	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

- Data transmission becomes slow or interrupted. Data may be lost or incomplete.
- Task scheduling is affected. For example, Yarn tasks cannot be executed properly or fail to be executed due to timeout.
- Data processing is affected. For example, HDFS data synchronization fails or the data is inaccurate.
- System performance deteriorates. The efficiency and quality of data processing is low.

### Possible Causes

- The node breaks down.
- The network is faulty.

### Procedure

**Check the network health status.**

**Step 1** In the alarm list on FusionInsight Manager, click the drop-down button of the alarm and view **Additional Information**. Record the source IP address and destination IP address of the node for which the alarm is reported.

**Step 2** Log in to the node for which the alarm is reported. On the node, ping the target node to check whether the network between the two nodes is normal.

- If yes, go to [6](#).
- If no, go to [3](#).

**Check the node status.**

**Step 3** On FusionInsight Manager, click **Host** and check whether the host list contains the faulty node to determine whether the faulty node has been removed from the cluster.

- If yes, go to [5](#).
- If no, go to [4](#).

**Step 4** Check whether the faulty node is powered off.

- If yes, start the faulty node and go to [Step 2](#).
- If no, contact related personnel to find root cause, if need to remove the faulty nodes from the cluster and go to [5](#), otherwise go to [6](#).

**Step 5** Remove the file `$NODE_AGENT_HOME/etc/agent/hosts.ini` of all nodes in the cluster, and clean up the file `/var/log/Bigdata/unreachable/unreachable_ip_info.log`, and then manually clear the alarm.


**Step 6** Wait for 30 seconds and checking if the alarm was been cleared.

- If yes, no further action is required.
- If no, go to [7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **OmmAgent** from the **Service** and click **OK**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.63 ALM-12091 Abnormal disaster Resources

### Alarm Description

HA checks the disaster resources of Manager every 86 seconds. This alarm is generated when HA detects that the disaster resources have been abnormal for 10 consecutive times.

This alarm is cleared when HA detects that the disaster resources become normal.

**Resource Type** of disaster is **Single-active**. Active/Standby switchover will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new disaster resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12091	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System



- The active/standby Manager switchover occurs.
- The disaster process restarts repeatedly, which may cause active/standby DR to be unavailable.

### Possible Causes

The disaster process is abnormal.

### Handling Procedure

**Check whether the disaster process is normal.**

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click  to view the name of the host for which the alarm is generated.
- Step 2** Log in to the host for which the alarm is generated as user **root**.
- Step 3** Run the **su - omm** command to switch to user **omm**.
- Step 4** Run the **sh \${BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** command to check whether the status of the disaster resources managed by the HA is normal. In the single-node system, the disaster resource is in the normal state. In the dual-node system, the disaster resource is in the normal state on the active node and in the stopped state on the standby node.
- If yes, go to **Step 7**.
  - If no, go to **Step 5**.
- Step 5** Run the **vi \${BIGDATA\_LOG\_HOME}/disaster/disaster.log** command to check whether the disaster resource log of HA contains the keyword **ERROR**. If yes, analyze the logs to locate the resource exception cause and fix the exception.
- Step 6** Wait 5 minutes and check whether the alarm is automatically cleared.
- If yes, no further action is required.
  - If no, go to **Step 7**.
- Collect fault information.**
- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 8** Expand the **Service** drop-down list, select **Disaster** for the target cluster, and click **OK**.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.64 ALM-12099 core dump Occurred

### Description

GaussDB A manages the core file. It manages the lifecycle of core files generated when applications crash and manages alarm notification. This alarm is generated when a new core file is detected.

 NOTE

This section applies to MRS 3.1.5 and later versions.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
12099	Minor	No

## Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

If a key process crashes, the cluster may be unavailable for a short period of time.

## Possible Causes

Related processes crash.

## Procedure

---

 CAUTION

- Users' sensitive data may be involved in the following operations on parsing the core file stack information. Therefore, development or O&M personnel can perform these operations only after being authorized.
  - The core file generated for the alarm is retained for 72 hours by default. If the period for storing the file exceeds 72 hours or the file size exceeds the preset value, the system automatically clears the file. Therefore, once this alarm is generated, contact O&M personnel in a timely manner.
- 

**Step 1** In the alarm list on the FusionInsight Manager page, click the row containing the alarm, and view the host IP address for which the alarm is generated in the alarm

details. Then, view the path for storing the core file according to the **DumpedFilePath** attribute in the additional information.

**Step 2** Log in to the host for which the alarm is generated as user **omm** and run the **gdb --version** command to check whether the gdb tool is installed on the host.

- If no, install the gdb tool and then go to [Step 3](#).
- If yes, go to the [Step 3](#).

**Step 3** Use the gdb tool to view the stack details of the core file.

1. Go to the **DumpedFilePath** directory and find the core file.
2. Run the following commands to obtain the symbol table of the core file:

```
source $BIGDATA_HOME/mppdb/.mppdbgs_profile
```

```
cd ${BIGDATA_HOME}/FusionInsight_MPPDB_XXX/install/FusionInsight-MPPDB-XXX/package/MPPDB_ALL_PACKAGE
```

```
tar -xzf GaussDB-Kernel-V300R002C00-Operating system-64bit-symbol.tar.gz
```

```
cd symbols/bin/
```

Find the symbol table file whose name is the same as the process name in the alarm. For example, the symbol table file for the **cm\_agent** process is **cm\_agent.symbol**.

Copy the obtained symbol table to the **`\${GAUSSHOME}/bin** directory.

3. Run the **gdb --batch -n -ex thread -ex bt core file name** command to view the stack details of the core file.

**Step 4** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system does not automatically clear this alarm, and you need to manually clear the alarm.

## Related Information

None

## 7.12.65 ALM-12100 AD Service Connection Failed

### Alarm Description

After a third-party active directory (AD) is interconnected, the third-party AD domain user can be synchronized using the synchronization period (60 minutes by default) or manually. During data synchronization, the AD service status will be checked. This alarm is generated when AD service unavailability is detected for three consecutive times. This alarm is cleared when AD service recovers.

#### NOTE

This section applies to MRS 3.1.5 and later versions.

## Alarm Attributes

Alarm ID	Severity	Auto Clear
12100	Major	Yes

## Alarm Parameters

Alarm Parameters	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.
RoleName	Specifies the name of the role for which the alarm is generated.
HostName	Specifies the name of the host for which the alarm is generated.

## Impact on the System

When the alarm is generated, the AD service is unavailable, and AD domain user synchronization fails. An AD domain user cannot log in to FusionInsight Manager and execute services.

## Possible Causes

- The configuration item for interconnecting with the third-party AD is incorrect.
- The network connection between FusionInsight and the third-party AD service is faulty.
- AD server fault
- AD service fault

## Handling Procedure

**Check the third-party AD configuration.**

**Step 1** On the **FusionInsight Manager** page, choose **System > Permission > Third-Party AD**. The third-party AD configuration page is displayed.

**Step 2** Check whether the **AD IP Address**, **LDAP Port**, **Bind DN**, and **Bind DN Password** parameters are correctly set.

- If yes, go to [Step 5](#).
- If no, go to [Step 3](#).

**Step 3** Modify the incorrect parameters, and then click **OK**.

**Step 4** Choose **System > Permission > User > AD Domain User**, click **Manual Synchronization**, and check whether the message "Manual synchronization successfully." is displayed in the upper right corner of the page.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Check the third-party AD server and the network.**

**Step 5** Log in to the active management node as user **root**.

**Step 6** On the host you have logged in to, ping the IP address of the third-party AD server to check whether the third-party AD server can be pinged.

- If yes, go to [Step 7](#).
- If no, go to [Step 8](#).

**Step 7** Run the following command to check whether the third-party AD service can be connected:

```
telnet IP port
```

*IP* indicates the IP address of the third-party AD server, and *port* indicates the port used by the third-party AD server.

- If yes, go to [Step 8](#).
- If no, contact O&M personnel to check the network.


**Step 8** Contact the third-party AD service administrator to check whether the AD service is normal.

- If yes, go to [Step 9](#).
- If no, contact the third-party AD service administrator to rectify the AD server fault.

**Collect the fault information.**

**Step 9** On the **FusionInsight Manager** page, choose **O&M > Log > Download**.

**Step 10** In the **Service** area, select **Controller** under **OMS** and click **OK**.

**Step 11** Click  in the upper right corner to set **Start Date** and **End Date** to 10 minutes before and after the time when the alarm is generated, and click **Download**.

**Step 12** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm. You do not need to manually clear it.

## Reference

None



## 7.12.66 ALM-12101 AZ Unhealthy

### Description

After the AZ DR function is enabled, the system checks the AZ health status every 5 minutes. This alarm is generated when the system detects that the AZ is subhealthy or unhealthy. This alarm is cleared when the AZ becomes healthy.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
12101	Critical	Yes

### Parameters

Parameter	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
AZName	Specifies the AZ for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

The health status of an AZ is determined by whether the health status of storage resources (HDFS), computing resources (Yarn), and key roles in the AZ exceeds the configured threshold.

An AZ is subhealthy when:

- The computing resources (Yarn) are unhealthy, but the storage resources (HDFS) are healthy. Tasks cannot be submitted to the local AZ, but data can still be read and written in the local AZ.
- The computing resources (Yarn) are healthy, but some storage resources (HDFS) are unhealthy. Tasks can be submitted to the local AZ, and some data can be read and written in the local AZ. This depends on the locality of data detected by Spark/Hive scheduling.

An AZ is unhealthy when:

- The computing resources (Yarn) are healthy, but the storage resources (HDFS) are unhealthy. Although tasks can be submitted to the local AZ, data cannot be read or written in the local AZ. As a result, the tasks submitted to the local AZ are invalid.

- The computing resources (Yarn) and storage resources (HDFS) are unhealthy. Tasks cannot be submitted to the local AZ, and data cannot be read or written in the local AZ.
- The health status of key roles except Yarn and HDFS is lower than the configured threshold.

## Possible Causes

- The computing resources (Yarn) are unhealthy.
- The storage resources (HDFS) are unhealthy.
- Some storage resources (HDFS) are unhealthy.
- Key roles except Yarn and HDFS are unhealthy.

## Procedure

### Disable the DR drill.

- Step 1** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Cross-AZ HA**. The Cross-AZ HA page is displayed.
- Step 2** In the AZ DR list, check whether **Perform DR Drill** in the **Operation** column of the AZ whose health status is **Unhealthy** is gray.
- If yes, go to **Step 4**.
  - If no, go to **Step 3**.
- Step 3** Click **Restore** in the **Operation** column of the target AZ. Wait 2 minutes and refresh the page to view the health status of the AZ. Check whether the health status is normal.
- If yes, no further action is required.
  - If no, go to **Step 4**.

### Collect the fault information.

- Step 4** Log in to the active management node as user **root**.
- Step 5** View logs of unhealthy services.
- HDFS log files are stored in **/var/log/Bigdata/hdfs/nn/hdfs-az-state.log**.
  - Yarn log files are stored in **/var/log/Bigdata/yarn/rm/yarn-az-state.log**.
  - For other services, view the service health check logs in the corresponding service log directory.
- Step 6** Contact O&M personnel and provide detailed log file information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.67 ALM-12102 AZ HA Component Is Not Deployed Based on DR Requirements

### Alarm Description

The alarm module checks the deployment status of AZ HA components every 5 minutes. This alarm is generated when the components that support DR are not deployed based on DR requirements after AZ is enabled. This alarm is cleared when the components are deployed based on DR requirements.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12102	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

### Impact on the System

The cross-AZ HA capability of a single cluster is affected.


### Possible Causes

The roles of the components that support DR are not deployed based on DR requirements.

### Handling Procedure

**Obtain alarm information.**

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

**Step 2** In the alarm list, click  in the row that contains the alarm and view the roles that are not deployed based on DR requirements in **Additional Information**.

**Redeploy the role instance.**

**Step 3** Choose **Cluster > Services > Name of the desired service > Instance**. On the instance page, redeploy or adjust the role instance.

**Step 4** Check whether the alarm is cleared 10 minutes later.

- If yes, no further action is required.
- If no, contact O&M personnel.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.68 ALM-12103 Executor Resource Exception

### Alarm Description

HA checks the Executor resources of Manager every 30 seconds. This alarm is generated when HA detects that the Executor resources are abnormal for two consecutive times.

This alarm is cleared when the Executor resources are normal.

**Resource Type** of Executor is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new Executor resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12103	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System


- The active/standby Manager switchover occurs.
- The Executor process keeps restarting. As a result, the cluster page may fail to be accessed.

## Possible Causes

The Executor process is abnormal.

## Handling Procedure

**Check whether the Executor process is abnormal.**

**Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click  to view the name of the host for which the alarm is generated.

**Step 2** Log in to the host for which the alarm is generated as user **root**.

**Step 3** Run the **su - omm** command to switch to user **omm**.

**Step 4** Run the **sh \${BIGDATA\_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status\_ha.sh** command to check whether the status of the Executor resources managed by the HA is normal. In the single-node system, the Executor resource is in the normal state. In the dual-node system, the Executor resource is in the normal state on the active node and in the stopped state on the standby node.

- If yes, go to [Step 7](#).
- If no, go to [Step 5](#).

**Step 5** Run the **vi \$BIGDATA\_LOG\_HOME/omm/oms/ha/scriptlog/executor.log** command to check whether the Executor resource log of HA contains the keyword **ERROR**. If yes, analyze the log to locate the resource exception cause and fix the exception.


**Step 6** After 5 minutes, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect the fault information.**

**Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 8** In the **Services** area, select **Controller** and **OmmServer**, and click **OK**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact O&M personnel and provide the collected logs.

----End

## 7.12.69 ALM-12104 Abnormal Knox Resources

### Alarm Description

HA checks the Knox resources of Manager every 70 seconds. This alarm is generated when HA detects that the Knox resources are abnormal for three consecutive times.

This alarm is cleared when HA detects that the Knox resources are normal.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12104	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System

Requests sent by upper-layer services by using Knox cannot be properly processed.


### Possible Causes

The Knox process is abnormal.

### Handling Procedure

Check whether the Knox process is normal.

- Step 1** Log in to FusionInsight Manager. In the alarm list, locate the row that contains the alarm and view the name of the host for which the alarm is generated.
- Step 2** Use PuTTY to log in to the host for which the alarm is generated as user **root**.
- Step 3** Run the **su - omm** command to switch to user **omm**.

- Step 4** Run the `sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh` command to check whether the status of the Knox resources managed by HA is normal. If the status is normal, the Knox resources are normal. Otherwise, the Knox resources are abnormal.
- If yes, go to [Step 7](#).
  - If no, go to [Step 5](#).
- Step 5** Run the `vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/knox.log` command to check whether the Knox resource log of HA contains the keyword **ERROR**. If yes, analyze the log to locate the resource exception cause and fix the exception.
- Step 6** After 5 minutes, check whether this alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 7](#).
- Collect the fault information.
- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 8** In the **Services** area, select **Controller** and **OmmServer**, and click **OK**.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.
- End

## 7.12.70 ALM-12110 Failed to get ECS temporary AK/SK

### Alarm Description

Meta calls the ECS API to obtain the AK/SK information every 5 minutes and caches the information. Before the AK/SK expires, Meta calls the API again to update it. This alarm is generated when Meta fails to call the API for three consecutive times.

This alarm is cleared when Meta successfully calls the ECS API.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12110	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System


The cluster cannot obtain the latest temporary AK/SK. In the storage and compute separation scenario, OBS may fail to be accessed. As a result, component services cannot be properly processed.

## Possible Causes

- The meta role of the MRS cluster is abnormal.
- The cluster has been bound to an agency and accessed OBS but has been unbound from the agency. As a result, the cluster has not been bound to any agency.

## Handling Procedure

**Check the status of the meta role.**

- Step 1** On FusionInsight Manager of the cluster, choose **O&M > Alarm > Alarms**. On the page that is displayed, click  in the row containing the alarm, and determine the IP address of the host for which the alarm is generated.
- Step 2** On FusionInsight Manager of the cluster, choose **Cluster > Services > meta**. On the page that is displayed, click the **Instances** tab, and check whether the meta role corresponding to the host for which the alarm is generated is normal.
  - If yes, go to [Step 5](#).
  - If no, go to [Step 3](#).
- Step 3** Select the abnormal role, click **More**, and select **Restart Instance** to restart the abnormal meta role.
- Step 4** Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 5](#).
- Step 5** Log in to the host obtained in [Step 1](#) and check whether the `/var/log/Bigdata/meta/mrs-meta.log` file contains error information. If yes, rectify the fault based on the log information.



**Step 6** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Rebind the cluster to an agency.**

**Step 7** Log in to the MRS management console.

**Step 8** In the navigation pane on the left, choose **Active Clusters**. On the page that is displayed, click the cluster name to go to its overview page. Then, check whether the cluster is bound to an agency in the O&M management area.

- If yes, go to [Step 10](#).
- If no, go to [Step 9](#).


**Step 9** Click **Manage Agency**. On the page that is displayed, rebind the cluster to an agency. Then check whether the alarm is cleared a few minutes later.

- If yes, no further action is required.
- If no, go to [Step 10](#).

**Collect fault information.**

**Step 10** On FusionInsight Manager of the active cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 11** Expand the **Service** drop-down list, select **meta** for the target cluster, and click **OK**.

**Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact O&M personnel and provide the collected logs.

----End

## 7.12.71 ALM-12172 Failed to Report Metrics to Cloud Eye

### Alarm Description

After metric sharing is enabled for a cluster, the Controller periodically collects cluster metrics and reports them to Cloud Eye.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12172	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

MRS monitoring metrics are unavailable on Cloud Eye.


## Possible Causes

- Failed to call Cloud Eye APIs due to insufficient permissions.
- Failed to report data to Cloud Eye due to network problems.
- Failed to report data to Cloud Eye due to internal errors.


## Handling Procedure

**Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the additional information in the alarm details.

**Step 2** Rectify the fault based on the following scenarios:

- If "Call CES to send metrics fail. Permission exception" is displayed in the additional information, the token of the resource tenant is invalid. Restart the Controller and obtain the token again. The detailed procedure is as follows:
  - a. Log in to FusionInsight Manager, click the task center icon  in the upper right corner, and verify that no task is being executed in the task center.
  - b. Use PuTTY to log in to the active management node as user **omm**.
  - c. Restart the controller.  
**sh \${BIGDATA\_HOME}/om-server/om/sbin/restart-controller.sh**
  - d. Wait for 5 minutes and check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to **Step 3**.
- If "Call CES to send metrics fail. Request CES error code xxx" or "CES internal error code xxx" is displayed, the function service is abnormal. Go to **Step 3**.
- If "Call CES to send metrics fail. Too many request" is displayed in additional alarm information, the service request triggers traffic control. Go to **Step 3**.

### Collect fault information.

- Step 3** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 4** On the displayed page, select **Controller** of **OMS** for **Service**, and click **OK**. Select the active node for **Host** and click **OK**.
- Step 5** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 6** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.72 ALM-12180 Suspended Disk I/O

### Alarm Description

**For MRS 3.3.0 and its later versions as well as MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions:**

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 6 seconds within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. The disk queue depth (**avgqu-sz**) is greater than 0, the IOPS or bandwidth is 0, and **ioutil** is greater than 99% in at least 10 collection periods within 30 seconds.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 1000 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 3 seconds within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. The disk queue depth (**avgqu-sz**) is greater than 0, the IOPS or bandwidth is 0, and **ioutil** is greater than 99% in at least 10 collection periods within 30 seconds.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 500 ms within 300 seconds.

The collection period is 3 seconds, and the detection period is 30 or 300 seconds. This alarm is automatically cleared when neither of the preceding conditions is met for three consecutive detection periods (30 or 300 seconds).

**For versions earlier than MRS 3.3.0:**

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency exceeds 6 seconds within 30 seconds in at least 10 collection periods.
  - By default, the system collects data every 3 seconds. The disk queue depth (**avgqu-sz**) is greater than 0, the IOPS or bandwidth is 0, and **ioutil** is greater than 99% in at least 10 collection periods within 30 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency exceeds 2 seconds within 30 seconds in at least 10 collection periods.
  - By default, the system collects data every 3 seconds. The disk queue depth (**avgqu-sz**) is greater than 0, the IOPS or bandwidth is 0, and **ioutil** is greater than 99% in at least 10 collection periods within 30 seconds.

This alarm is automatically cleared when the preceding conditions have not been met for 90s.

 **NOTE**

For details about how to obtain and calculate related parameters, see [Related Information](#).

**Alarm Attributes**

Alarm ID	Alarm Severity	Auto Cleared
12180	Major	Yes

**Alarm Parameters**

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
DiskName	Specifies the disk for which the alarm was generated.

## Impact on the System

If the I/O usage keeps increasing, operations will be affected and services will be interrupted. The possible impacts are as follows:

- The system I/O performance deteriorates, which means slow response and low throughput. For example, job submission is slow, page responds slowly, interface response times out, and the system is in error or even crash.
- Customer services may be interrupted. The system may break down and the key information stored on the faulty disk may be lost.

## Possible Causes

The disk is aged.

## Handling Procedure

**Replace the disk.**

**Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**.

**Step 2** View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is reported.

**Step 3** Replace the hard disk.


**Step 4** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Collect fault information.**

**Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** Select **OMS** for **Service** and click **OK**.

**Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

**Obtain and calculate related parameters as follows:**

- Run the following command in the OS to collect data:  
**iosat -x -t 1 1**

```
[root@node-master1hxyk ~]# iostat -x -t 1 1
Linux 4.18.0-147.5.2.10.el8.x86_64 (node-master1hxyk) 10/12/2022 _x86_64_ (8 CPU)

10/12/2022 05:24:09 PM
avg-cpu: %user %nice %system %iowait %steal %idle
 24.49 0.00 13.82 0.11 0.00 61.58

Device: r/s kB/s rrqm/s wrqm/s r_wait rreqsz w/s kB/s wrqm/s w_wait wreqsz d/s kB/s drqm/s wrqm d_wait dreqsz aqu-sz %util
da-0 0.00
da-1 0.07 0.30 0.00 0.00 0.00 0.00 4.41 0.90 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.03 0.61
vda 1.90 61.59 0.02 0.96 1.65 32.43 22.16 403.26 33.50 60.19 1.80 18.20 0.00 0.00 0.00 0.00 0.00 0.00 0.03 1.80
vdb 0.11 2.51 0.00 0.01 0.08 22.22 24.05 351.18 15.74 41.03 1.02 14.60 0.00 0.00 0.00 0.00 0.00 0.01 1.59
```

Parameters are as follows:

**avgqu-sz** indicates the disk queue depth.

The sum of **r/s** and **w/s** is the IOPS.

The sum of **rkB/s** and **wkB/s** is the bandwidth.

**%util** is the **ioutil** value.

- MRS 3.1.0:

Run the **iostat -x -t** command in the OS.

```
[omm@node-master1hxyk ~]# iostat -x -t
Linux 3.10.0-862.14.1.5.el7.elrepo.x86_64 (node-master1hxyk) 11/11/2022 _x86_64_ (4 CPU)

11/11/2022 03:35:20 PM
avg-cpu: %user %nice %system %iowait %steal %idle
 27.66 0.00 15.66 0.63 0.00 56.06

Device: rrqm/s wrqm/s r/s w/s kB/s kB/s avgrq-sz avgqu-sz await r_await w_await svctm %util
vda 0.13 29.26 1.71 23.51 187.56 608.08 63.11 0.91 36.02 50.86 34.94 0.64 1.62
vdb 0.00 14.45 0.08 27.34 1.35 301.81 22.12 0.08 2.81 26.57 2.74 0.53 1.45
```

- Calculate **svctm** as follows in versions later than MRS 3.1.0:  

$$svctm = (tot\_ticks\_new - tot\_ticks\_old) / (rd\_ios\_new + wr\_ios\_new - rd\_ios\_old - wr\_ios\_old)$$
- Versions earlier than MRS 3.3.0: If  $rd\_ios\_new + wr\_ios\_new - rd\_ios\_old - wr\_ios\_old = 0$ , then **svctm** = 0.
- MRS 3.3.0 and its later versions as well as MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions:

When the detection period is 30 seconds, if  $rd\_ios\_new + wr\_ios\_new - rd\_ios\_old - wr\_ios\_old = 0$ , then **svctm** = 0.

When the detection period is 300 seconds and  $rd\_ios\_new + wr\_ios\_new - rd\_ios\_old - wr\_ios\_old = 0$ , if  $tot\_ticks\_new - tot\_ticks\_old = 0$ , then **svctm** = 0; otherwise, the value of **svctm** is infinite.

The parameters can be obtained as follows:

The system runs the **cat /proc/diskstats** command every 3 seconds to collect data. For example:

```
[omm@ ~]# cat /proc/diskstats
253 0 vda 1101553 35446 83439787 3338546 28744856 48314024 1054257852 52667332 0 19566526 40342913 0 0 0 0
253 1 vda1 596976 25494 54533791 2565698 5446004 8749340 215777628 121114542 0 6475805 11339691 0 0 0 0
253 2 vda2 15 0 108 136 0 0 0 0 150 129 0 0 0 0
253 5 vda5 22373 1364 2525122 79502 4212374 4104759 161597984 8145606 0 3598808 6239095 0 0 0 0
253 6 vda6 11145 314 529002 85050 259201 70368 4412408 321454 0 189336 259725 0 0 0 0
253 7 vda7 157987 105 3477434 149542 6507077 1028968 140666992 14349866 0 1679035 11116587 0 0 0 0
253 8 vda8 312935 8169 22369722 458354 12179958 34360589 531802640 17724858 0 9060731 11385470 0 0 0 0
253 16 vdb 275920 21939 15977738 2171665 39472291 28236575 2653825040 482230505 0 30580346 465962048 0 0 0 0
253 17 vdb1 275439 21939 15948866 2171472 31290400 28236555 2653824832 481837775 0 30036724 465855080 0 0 0 0
7 0 loop0 356 0 17442 150 0 0 0 0 149 105 0 0 0 0

[omm@ ~]# cat /proc/diskstats
253 0 vda 1101553 35446 83439787 3338546 28747977 48319338 1054352084 52672715 0 19571460 40346640 0 0 0 0
253 1 vda1 596976 25494 54533791 2565698 5446015 8750402 215791076 12115169 0 6474429 11339985 0 0 0 0
253 2 vda2 15 0 108 136 0 0 0 0 150 129 0 0 0 0
253 5 vda5 22373 1364 2525122 79502 4105244 161614088 8146153 0 3599216 6239432 0 0 0 0
253 6 vda6 11145 314 529002 85050 259245 70433 4413368 321489 0 189389 259730 0 0 0 0
253 7 vda7 157987 105 3477434 149542 6507759 1029060 14067782 14351373 0 1679157 11117724 0 0 0 0
253 8 vda8 312935 8169 22369722 458354 12181277 34364199 531855680 17727525 0 9061647 11387424 0 0 0 0
253 16 vdb 275920 21939 15977738 2171665 39477604 28238831 2653881640 482234435 0 30581946 465964144 0 0 0 0
253 17 vdb1 275439 21939 15948866 2171472 31293358 28238811 2653881432 481841639 0 30038274 465857164 0 0 0 0
7 0 loop0 356 0 17442 150 0 0 0 0 149 105 0 0 0 0
```

In the two commands:

In the data collected for the first time, the number in the fourth column is the **rd\_ios\_old** value, the number in the eighth column is the **wr\_ios\_old** value, and the number in the thirteenth column is the **tot\_ticks\_old** value.

In the data collected for the second time, the number in the fourth column is the **rd\_ios\_new** value, the number in the eighth column is the **wr\_ios\_new** value, and the number in the thirteenth column is the **tot\_ticks\_new** value.

In this case, the value of **svctm** is as follows:

$$(19571460 - 19569526) / (1101553 + 28747977 - 1101553 - 28744856) = 0.6197$$

## 7.12.73 ALM-12186 CGroup Task Usage Exceeds the Threshold

### Alarm Description

The system checks the CGroup task usage of user **omm** every 5 minutes. This alarm is generated when the CGroup task usage exceeds 90%. This alarm is cleared when the CGroup task usage is less than or equal to 90%.

CGroup task usage = Number of used CGroup tasks/Maximum number of CGroup tasks

You can run the **systemctl status user-\$(id -u).slice | grep limit | awk -F ' '{print \$2}'** command as user **omm** to obtain the number of used CGroup tasks of this user and run the **echo \$(systemctl status user-\$(id -u).slice | grep limit | awk -F ' '{print \$4}' | sed -e 's/)//g'** command to obtain the maximum number of CGroup tasks allowed for this user.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12186	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

- Failed to switch to user **omm**.
- Failed to create new **omm** processes.


- A faulty service or process cannot be restarted.

## Possible Causes

The CGroup task usage exceeds 90%.

## Handling Procedure

**Check the maximum number of threads that can be concurrently opened by user omm is properly set.**

**Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. On the page that is displayed, click  in the row containing the alarm, and view the name of the host for which the alarm is generated in **Location**. Click the host name to view its IP address.

**Step 2** Log in to the host for which the alarm is generated as user **omm**.

**Step 3** Run the following command to obtain the maximum number of threads that can be concurrently opened by user **omm** and check whether this number is greater than or equal to **60000**:

```
systemctl status user-$(id -u).slice | grep limit
```

- If yes, go to [Step 6](#).
- If no, go to [Step 4](#).

**Step 4** Switch to user **root** and run the following command to change the value for user **omm** to **60000**:

```
systemctl set-property user-2000.slice TasksMax=60000
```


**Step 5** Change the value of **UserTasksMax** in the **/etc/systemd/logind.conf** file to **60000**. (If the parameter is commented out, uncomment it.) Save the file, wait 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect fault information.**

**Step 6** On FusionInsight Manager of the cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, select **OmmServer** and **NodeAgent** for the target cluster, and click **OK**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.



## Related Information

None.

## 7.12.74 ALM-12187 Failed to Expand Disk Partition Capacity

### Alarm Description

The system checks the disk space every 60 seconds. When detecting that the disk space is expanded, the system expands disk partition. This alarm is generated when the disk partition fails to be expanded.

This alarm is cleared when the system detects that the disk partition is successfully expanded.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12187	Minor	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
MountDirectoryName	Specifies the directory for which the alarm is generated.

### Impact on the System

A disk partition scale-out failure may have the following impacts on the system:

- Data may be lost. (It is important to back up key information before scale-out.)
- The system may be unstable or cannot be started because the system files may be damaged.
- The disk may be unavailable. In this case, you need to format the disk and partition it again.

- The system performance may deteriorate because the disk partition space is not enough. The performance remains low until the disk partition is scaled out successfully.

## Possible Causes

- The growpart scale-out tool is not installed.
- The system fails to execute the command for expanding disk partition.

## Handling Procedure

### Check whether growpart is installed.

**Step 1** Log in to FusionInsight Manager, click **O&M**, and choose **Alarm > Alarms** to view the alarm details. In the **Location** column, check the name of the host and mount directory for which the alarm is generated. Click the host name to view its IP address.

**Step 2** Log in to the node for which the alarm is generated as user **root**.

**Step 3** Run the following command to check whether growpart is installed:

### which growpart

If information similar to the following is displayed, the growpart tool is installed. Otherwise, contact O&M personnel to install the growpart tool.

```
[root@xxx ~]#which growpart
/usr/bin/growpart
```

**Step 4** Wait for 5 minutes, then choose **O&M**, and choose **Alarm > Alarms** on FusionInsight Manager. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

### Run the disk partition expansion command.

**Step 5** Run the following command to view the disk and partition information:

### lsblk

Search for the partition and the disk based on the mount directory name in the alarm location information, and check the disk and partition sizes.

In the following example, the mount directory is **/srv/BigData/data1**, the used disk is **/dev/vdb**, and the disk partition is **/dev/vdb1**.

```
[root@xxx ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
loop0 7:0 0 1.5G 0 loop /media
vda 253:0 0 500G 0 disk
├─vda1 253:1 0 220G 0 part /
├─vda2 253:2 0 1K 0 part
├─vda5 253:5 0 10G 0 part /tmp
├─vda6 253:6 0 10G 0 part /var
├─vda7 253:7 0 60G 0 part /srv/BigData
└─vda8 253:8 0 180G 0 part /var/log
vdb 253:16 0 650G 0 disk
└─vdb1 253:17 0 600G 0 part /srv/BigData/data1
```

**Step 6** Run the following command to expand the partition using growpart:

```
growpart Data disk Partition number
```

Run the following command:

```
growpart /dev/vdb 1
```

If information similar to the following is displayed, the execution is successful. If the execution fails, contact O&M personnel.

```
[root@host- ~]# growpart /dev/vdb 1
CHANGED: partition=1 start=2048 old: size=1258287104 end=1258289152 new: size=1363146719 end=1363148767
```

**Step 7** Run the following command to expand the file system size of the disk partition:

```
resize2fs Disk partition
```

Run the following command:

```
resize2fs /dev/vdb1
```

If information similar to the following is displayed, the execution is successful:

```
[root@host- ~]# resize2fs /dev/vdb1
resize2fs 1.46.4 (18-Aug-2021)
Filesystem at /dev/vdb1 is mounted on /srv/BigData/data1; on-line resizing required
old_desc_blocks = 75, new_desc_blocks = 82
The filesystem on /dev/vdb1 is now 170393339 (4k) blocks long.
```

**Step 8** Wait for 5 minutes, click **O&M**, and choose **Alarm > Alarms** on FusionInsight Manager. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, contact O&M personnel.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.75 ALM-12188 diskmgmt Disk Monitoring Unavailable

### Alarm Description

NodeAgent checks the status of the diskmgmt disk monitoring service every 5 minutes. This alarm is generated when diskmgmt disk monitoring is unavailable.

This alarm is cleared when the diskmgmt disk monitoring service recovers.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12188	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

When diskmgmt disk monitoring is unavailable, the read-only detection of the device partition file system, device partition loss detection, and disk partition scale-out detection cannot be performed.

## Possible Causes

- The diskmgmt disk monitoring service does not exist.
- The diskmgmt disk monitoring service is not started.

## Handling Procedure

**Check whether the diskmgmt disk monitoring service exists.**

**Step 1** Log in to FusionInsight Manager, click **O&M**, and choose **Alarm > Alarms** to view the alarm details. In the **Location** column, check the name of the host for which the alarm is generated. Click the host name to view its IP address.

**Step 2** Log in to the node for which the alarm is generated as user **root**.

**Step 3** Run the following command to check whether the core service file exists:

```
stat /usr/local/diskmgmt/inner/diskmgtd
```

If the file does not exist, contact O&M personnel.

**Start the diskmgmt disk monitoring service.**

**Step 4** Run the following command to start the diskmgmt disk monitoring service:

```
systemctl restart diskmgmt
```

**Step 5** Run the following command to check whether the diskmgmt disk monitoring service is started:

#### systemctl status diskmgmt

- If information similar to the following is displayed, the service is started successfully. Go to [Step 6](#).

```
[root@host ~]# systemctl status diskmgmt
● diskmgmt.service - Disk monitor service
 Loaded: loaded (/usr/lib/systemd/system/diskmgmt.service; enabled; vendor preset: disabled)
 Active: active (running) since Thu 2023-08-10 09:29:18 CST; 5 days ago
 Main PID: 33996 (diskmgtd)
 Tasks: 2 (limit: 407663)
 Memory: 5.6M
 CGroup: /system.slice/diskmgmt.service
 └─ 33996 /bin/bash /usr/local/diskmgmt/inner/diskmgtd
 └─ 1974506 sleep 8

Notice: journal has been rotated since unit was started, output may be incomplete.
```

- If no, contact O&M personnel.

**Step 6** Wait for 5 minutes, click **O&M**, and choose **Alarm > Alarms** on FusionInsight Manager. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, contact O&M personnel.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.76 ALM-12190 Number of Knox Connections Exceeds the Threshold

### Alarm Description

The system periodically checks the number of connections to all Knox topologies. This alarm is generated when the number of connections to a topology exceeds the threshold (90% by default). This alarm is automatically cleared when the number of connections to a topology falls below the threshold.

#### NOTE

This alarm applies to clusters of MRS 3.1.0 or later.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12190	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Topology	Specifies the Knox topology for which the alarm was generated.

## Impact on the System

The topology may reach the upper limit of connections and fail to forward requests, adversely affecting the MRS functions.

## Possible Causes

Hue or Manager is too frequently used, but the default maximum number of Knox connections is small.

## Handling Procedure

**Step 1** Log in to active and standby OMS nodes as user **root**, respectively.

**Step 2** Add the following configuration to the **gateway-site.xml** file on the active and standby OMS nodes to increase the number of thread pools:

**vi /opt/knox/conf/gateway-site.xml**

```
<property>
<name>gateway.httpClient.maxConnections</name>
<value>64</value>
</property>
```

**Step 3** Log in to the active OMS node as user **omm** and run the following command to restart the Knox process:

**sh /opt/knox/bin/restart-knox.sh**

**Step 4** After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Step 5** Contact O&M personnel to rectify the fault.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.77 ALM-13000 ZooKeeper Service Unavailable

### Description

The system checks the ZooKeeper service status every 60 seconds. This alarm is generated when the ZooKeeper service is unavailable.

This alarm is cleared when the ZooKeeper service recovers.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
13000	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

ZooKeeper cannot provide coordination services for upper layer components and the components (such as Yarn and Flink) that depend on ZooKeeper may not run properly.

### Possible Causes

- The DNS is installed on the ZooKeeper node.
- The network is faulty.
- The KrbServer service is abnormal.

- The ZooKeeper instance is abnormal.
- The disk capacity is insufficient.

## Procedure

### Check the DNS.

- Step 1** Check whether the DNS is installed on the node where the ZooKeeper instance is located. On the Linux node where the ZooKeeper instance is located, run the `cat /etc/resolv.conf` command to check whether the file is empty.
- If yes, go to [Step 2](#).
  - If no, go to [Step 3](#).
- Step 2** Run the `service named status` command to check whether the DNS is started.
- If yes, go to [Step 3](#).
  - If no, go to [Step 5](#).
- Step 3** Run the `service named stop` command to stop the DNS service. If "Shutting down name server BIND waiting for named to shut down (28s)" is displayed, the DNS service is stopped successfully. Comment out the content (if any) in `/etc/resolv.conf`.
- Step 4** On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 5](#).

### Check the network status.

- Step 5** On the Linux node where the ZooKeeper instance is located, run the `ping` command to check whether the host names of other nodes where the ZooKeeper instance is located can be pinged successfully.
- If yes, go to [Step 9](#).
  - If no, go to [Step 6](#).
- Step 6** Modify the IP addresses in `/etc/hosts` and add the host name and IP address mapping.
- Step 7** Run the `ping` command again to check whether the host names of other nodes where the ZooKeeper instance is located can be pinged successfully.
- If yes, go to [Step 8](#).
  - If no, go to [Step 23](#).
- Step 8** On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 9](#).

### Check the KrbServer service status (Skip this step if the normal mode is used).

- Step 9** On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services**.
- Step 10** Check whether the KrbServer service is normal.



- If yes, go to [Step 13](#).
- If no, go to [Step 11](#).

**Step 11** Perform operations based on "ALM-25500 KrbServer Service Unavailable" and check whether the KrbServer service is recovered.

- If yes, go to [Step 12](#).
- If no, go to [Step 23](#).

**Step 12** On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 13](#).

**Check the ZooKeeper service instance status.**

**Step 13** On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > ZooKeeper > quorumpeer**.

**Step 14** Check whether the ZooKeeper instances are normal.

- If yes, go to [Step 18](#).
- If no, go to [Step 15](#).

**Step 15** Select instances whose status is not good, and choose **More > Restart Instance**.

 **NOTE**

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.

**Step 16** Check whether the instance status is good after restart.

- If yes, go to [Step 17](#).
- If no, go to [Step 18](#).

**Step 17** On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 18](#).

**Check disk status.**

**Step 18** On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Service > ZooKeeper > quorumpeer**, and check the node host information of the ZooKeeper instance.

**Step 19** On FusionInsight Manager, click **Host**.

**Step 20** In the **Disk** column, check whether the disk space of each node where ZooKeeper instances are located is insufficient (disk usage exceeds 80%).


- If yes, go to [Step 21](#).
- If no, go to [Step 23](#).

**Step 21** Expand disk capacity. For details, see "ALM-12017 Insufficient Disk Capacity".

**Step 22** On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 23](#).

**Collect fault information.**

- Step 23** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
  - Step 24** Select the following nodes in the required cluster from the **Service**: (KrbServer logs do not need to be downloaded in normal mode.)
    - ZooKeeper
    - KrbServer
  - Step 25** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
  - Step 26** Contact the O&M personnel and send the collected log information.
- End

**Alarm Clearing**

After the fault is rectified, the system automatically clears this alarm.

**Related Information**

None

**7.12.78 ALM-13001 Available ZooKeeper Connections Are Insufficient**

**Description**

The system checks ZooKeeper connections every 60 seconds. This alarm is generated when the system detects that the number of used ZooKeeper instance connections exceeds the threshold (80% of the maximum connections).

When the **Trigger Count** is 1, this alarm is cleared when the number of used ZooKeeper instance connections is smaller than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the number of used ZooKeeper instance connections is smaller than or equal to 90% of the threshold.

**Attribute**

Alarm ID	Alarm Severity	Auto Clear
13001	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host name for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

Available ZooKeeper connections are insufficient. When the connection usage reaches 100%, the ZooKeeper cannot process external connections. As a result, upstream components (such as Yarn and Flink) cannot run properly.

## Possible Causes

The number of connections to the ZooKeeper node exceeds the threshold. Connection leakage occurs on some connection processes, or the maximum number of connections does not comply with the actual scenario.

## Procedure

### Check connection status.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the displayed interface, click the drop-down button of **Available ZooKeeper Connections Are Insufficient** and confirm the node IP address of the host for which the alarm is generated in the Location Information.
- Step 2** Obtain the PID of the ZooKeeper process. Log in to the node involved in this alarm as user **root** and run the **pgrep -f proc\_zookeeper** command.
- Step 3** Check whether the PID can be correctly obtained.
  - If yes, go to **Step 4**.
  - If no, go to **Step 15**.
- Step 4** Obtain all the IP addresses connected to the ZooKeeper instance and the number of connections and check 10 IP addresses with top connections. Run the following command based on the obtained PID: **lsof -i|grep \$pid | awk '{print \$9}' | cut -d : -f 2 | cut -d \>-f 2 | awk '{a[\$1]++} END {for(i in a){print i,a[i] | "sort -r -g -k 2"}}' | head -10**. (The PID obtained in the preceding step is used.)

- Step 5** Check whether node IP addresses and number of connections are successfully obtained.
- If yes, go to [Step 6](#).
  - If no, go to [Step 15](#).
- Step 6** Obtain the ID of the port connected to the process. Run the following command based on the obtained PID and IP address: `lsof -i|grep $pid | awk '{print $9}'|cut -d \> -f 2 |grep $IP | cut -d : -f 2`. (The PID and IP address obtained in the preceding step are used.)
- Step 7** Check whether the port ID is successfully obtained.
- If yes, go to [Step 8](#).
  - If no, go to [Step 15](#).
- Step 8** Obtain the ID of the connected process. Log in to each IP address and run the following command based on the obtained port ID: `lsof -i|grep $port`. (The port ID obtained in the preceding step is used.)
- Step 9** Check whether the process ID is successfully obtained.
- If yes, go to [Step 10](#).
  - If no, go to [Step 15](#).
- Step 10** Check whether connection leakage occurs on the process based on the obtained process ID.
- If yes, go to [Step 11](#).
  - If no, go to [Step 12](#).
- Step 11** Close the process where connection leakage occurs and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 12](#).
- Step 12** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Configurations** > **All Configurations** > **quorumpeer** > **Performance** and increase the value of **maxCnxns** as required.

**Figure 7-65** maxCnxns

Parameter	Value
maxClientCnxns	2000
maxCnxns	20000

- Step 13** Save the configuration and restart the ZooKeeper service.

**NOTE**


The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

- Step 14** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 15](#).

**Collect fault information.**

**Step 15** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 16** Select **ZooKeeper** in the required cluster from the **Service**:

**Step 17** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 18** Contact the O&M personnel and send the collected log information.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.79 ALM-13002 ZooKeeper Direct Memory Usage Exceeds the Threshold

### Description

The system checks the direct memory usage of the ZooKeeper service every 30 seconds. The alarm is generated when the direct memory usage of a ZooKeeper instance exceeds the threshold (80% of the maximum memory).

When the **Trigger Count** is 1, this alarm is cleared when the ZooKeeper Direct memory usage is less than the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the ZooKeeper Direct memory usage is less than 80% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
13002	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the available memory of ZooKeeper is insufficient, memory overflow may occur and services may break down. As a result, upstream services (such as HDFS and Yarn) fail to run.

## Possible Causes

The direct memory of the ZooKeeper instance is overused or the direct memory is inappropriately allocated.

## Procedure

### Check the direct memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the displayed interface, click the drop-down button of **ZooKeeper Direct Memory Usage Exceeds the Threshold**. Check the IP address of the instance that reports the alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > ZooKeeper > Instance > quorumpeer(the IP address checked)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > CPU and Memory**, and select **ZooKeeper Heap And Direct Buffer Resource Percentage**, click **OK**.
- Step 3** Check whether the used direct buffer memory of ZooKeeper reaches 80% of the maximum direct buffer memory specified for ZooKeeper.
  - If yes, go to [Step 4](#).
  - If no, go to [Step 8](#).
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > ZooKeeper > Configurations > All Configurations > quorumpeer > System** to check whether "-XX:MaxDirectMemorySize" exists in the **GC\_OPTS** parameter.
  - If yes, in the **GC\_OPTS** parameter, delete "-XX:MaxDirectMemorySize" and go to [Step 5](#).
  - If no, go to [Step 6](#).

**Step 5** Save the configuration and restart the ZooKeeper service.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

**Step 6** Check whether the **ALM-13004 ZooKeeper Heap Memory Usage Exceeds the Threshold** exists.

- If yes, handle the alarm by referring to **ALM-13004 ZooKeeper Heap Memory Usage Exceeds the Threshold**.
- If no, go to [Step 7](#).


**Step 7** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect fault information.**

**Step 8** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 9** Select **ZooKeeper** in the required cluster from the **Service**.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.80 ALM-13003 GC Duration of the ZooKeeper Process Exceeds the Threshold

### Alarm Description

The system checks the garbage collection (GC) duration of the ZooKeeper process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
13003	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

The ZooKeeper process may respond slowly. Services of upper-layer components (such as Yarn, Flink, and Spark) may fail.

## Possible Causes

The heap memory of the ZooKeeper process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure


**Check the GC duration.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the displayed page, click the drop-down list of **GC Duration of the ZooKeeper Process Exceeds the Threshold**. View the IP address of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > ZooKeeper > Instance > quorumpeer**. Click the drop-down list in the upper right corner of **Chart**, choose **Customize > GC**, select **ZooKeeper GC Duration per Minute**, and click **OK** to check the GC duration statistics of the ZooKeeper process collected every minute.



**Figure 7-66** ZooKeeper GC duration

### Customize Statistics



**Step 3** Check whether the GC duration of the ZooKeeper process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 8](#).

**Step 4** Check whether memory leakage occurs in the application.

**Step 5** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > ZooKeeper > Configurations > All Configurations > quorumpeer > System**. Increase the value of the **GC\_OPTS** parameter as required.

**NOTE**

Generally, **-Xmx** is twice of ZooKeeper data capacity. If the capacity of ZooKeeper reaches 2 GB, set **GC\_OPTS** as follows:

```
-Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=512M -XX:MetaspaceSize=64M -XX:MaxMetaspaceSize=64M -XX:CMSFullGCsBeforeCompaction=1
```

**Step 6** Save the configuration and restart the ZooKeeper service.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.


**Step 7** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect the fault information.**

**Step 8** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 9** Expand the **Service** drop-down list, and select **ZooKeeper** for the target cluster.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.81 ALM-13004 ZooKeeper Heap Memory Usage Exceeds the Threshold

## Description

The system checks the heap memory usage of the ZooKeeper service every 60 seconds. The alarm is generated when the heap memory usage of a ZooKeeper instance exceeds the threshold (95% of the maximum memory).

The alarm is cleared when the memory usage is less than the threshold.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
13004	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the available ZooKeeper heap memory is insufficient, memory overflow may cause service breakdown, and upstream components (such as Yarn, Flink, and Spark) may fail to run.

## Possible Causes

The heap memory of the ZooKeeper instance is overused or the heap memory is inappropriately allocated.

## Procedure

### Check heap memory usage.

- Step 1** On the FusionInsight Manager portal, On the displayed interface, click the drop-down button of **ZooKeeper Heap Memory Usage Exceeds the Threshold** and confirm the node IP address of the host for which the alarm is generated in the Location Information.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > ZooKeeper > Instance**, click **quorumpeer** in the **Role** column of the corresponding IP address. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > CPU and Memory**, and select **ZooKeeper Heap And Direct Buffer Resource Percentage**, click **OK**. Check the heap memory usage.
- Step 3** Check whether the used heap memory of ZooKeeper reaches 95% of the maximum heap memory specified for ZooKeeper.
  - If yes, go to [Step 4](#).
  - If no, go to [Step 7](#).

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > ZooKeeper > Configurations > All Configurations > quorumpeer > System**. Increase the value of **-Xmx** in **GC\_OPTS** as required. The details are as follows:

1. On the **Instance** tab, click **quorumpeer** in the **Role** column of the corresponding IP address. Choose **Customize > CPU and Memory** in the upper right corner, and select **ZooKeeper Heap And Direct Buffer Resource**, click **OK** to check the heap memory used by ZooKeeper.
2. Change the value of **-Xmx** in the **GC\_OPTS** parameter based on the actual heap memory usage. Generally, the value is twice the size of the ZooKeeper data volume. For example, if 2 GB ZooKeeper heap memory is used, the following configurations are recommended: **-Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=512M -XX:MetaspaceSize=64M -XX:MaxMetaspaceSize=64M -XX:CMSFullGCsBeforeCompaction=1**

**Step 5** Save the configuration and restart the ZooKeeper service.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.


**Step 6** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **ZooKeeper** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.82 ALM-13005 Failed to Set the Quota of Top Directories of ZooKeeper Components

### Description

The system sets quotas for each ZooKeeper top-level directory in the **customized.quota** configuration item and components every 5 hours. This alarm is generated when the system fails to set the quota for a directory.

This alarm is cleared when the setting succeeds after a failure.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13005	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
ServiceDirectory	Specifies the directory for which the alarm is generated.
Trigger Condition	Specifies the cause of the alarm.

### Impact on the System

Components can write a large amount of data to the top-level directory of ZooKeeper. As a result, services or services of upstream components (such as Yarn, Flink, and Spark) that depend on the top-level directory are abnormal.

### Possible Causes

The quota for the alarm directory is inappropriate.

### Procedure

**Check whether the quota for the alarm directory is appropriate.**

- Step 1** Log in to FusionInsight Manager, and choose **Cluster > Name of the desired cluster > Services > ZooKeeper**. On the displayed page, choose **Configurations > All Configurations > Quota**. Check whether the directory for which the alarm is reported and its quota exist in the **customized.quota** configuration item.

- If yes, go to [Step 5](#).
- If no, go to [Step 2](#).

**Step 2** Check whether the alarm directory for which the alarm is reported is in the following alarm list.

**Table 7-100** Component alarm directory

Component	Alarm Directory
Hbase	/hbase
Hive	/beelinesql
Yarn	/rmstore
Storm	/stormroot
Streaming	/storm
Kafka	/kafka

- If yes, go to [Step 3](#).
- If no, go to [Step 7](#).

**Step 3** View the component of the alarm directory in the table, open the corresponding service page, and choose **Configurations > All Configurations**. On the displayed page, search for **zk.quota** in the upper right corner. The search result is the quota of the alarm directory.

**Step 4** Check whether the quota of the alarm directory for which the alarm is reported is appropriate. The quota must be greater than or equal to the actual value, which can be obtained in **Trigger Condition**.

**Step 5** Modify the **services.quota** value as prompted and save the configuration.

**Step 6** After the time specified by **service.quotas.auto.check.cron.expression**, check whether the alarm is cleared.


The **service.quotas.auto.check.cron.expression** parameter indicates the scheduled expression used by ZooKeeper to set the directory quota. You can choose **Cluster > Services > ZooKeeper > Configurations > All Configurations** on FusionInsight Manager and set this parameter. The default value is `*/5 * * * *`, indicating 5 minutes.

- If it is, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **ZooKeeper** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.83 ALM-13006 Znode Number or Capacity Exceeds the Threshold

### Description

The system periodically detects the status of secondary Znode in the ZooKeeper service data directory every four hours. This alarm is generated when the number or capacity of secondary Znodes exceeds the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13006	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
ServiceDirectory	Specifies the directory for which the alarm is generated.
Trigger Condition	Specifies the cause of the alarm.

## Impact on the System

A large amount of data is written to the ZooKeeper data directory space. As a result, services of upstream components (such as Yarn, Flink, and Spark) that

depend on this directory are abnormal. For details, see the alarm location information.

## Possible Causes


A large amount of data is written to the ZooKeeper data directory. The threshold is not appropriate.

## Procedure

**Check whether a large amount of data is written to the directory for which the alarm is generated.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the displayed interface, click the drop-down button of **Znode Number or Capacity Exceeds the Threshold**. Confirm the Znode for which the alarm is generated in Location Information.
- Step 2** Log in to FusionInsight Manager, open the ZooKeeper service interface, and select **Resource**. In the table **Used Resources (By Second-Level Znode)**, check whether a large amount of data is written to the top-level Znode for which the alarm is reported.
  - If it is, go to [Step 3](#).
  - If it is not, go to [Step 4](#).

**Figure 7-67** Used Resources (By Second-Level Znode)



- Step 3** Log in to the ZooKeeper client and delete the data in the top-level Znode.
- Step 4** Log in to FusionInsight Manager and open the ZooKeeper service interface. On the **Resource** page, choose [By Znode quantity](#) in **Used Resources (By Second-Level Znode)**. **Threshold Configuration of By Znode quantity** is displayed. Click **Modify** under **Operation**. Increase the threshold by referring to the value of **max.Znode.count** by choosing **Cluster > Name of the desired cluster > Services > ZooKeeper > Configurations > All Configurations > Quota**.



**Figure 7-68** Modify Rule

### Modify Rule

\* Rule Name:


\* Alarm Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Others

Thresholds:

Start and End Time	Threshold
<input type="text" value="00:00"/> - <input type="text" value="23:59"/>	<input type="text" value="200000"/>

**Step 5** In the **Used Resources (By Second-Level Znode)**, choose  > **By capacity**. The **Threshold Settings** page of **By Capacity** is displayed. Click **Modify** under **Operation**. Increase the threshold by referring to the value of **max.data.size** by choosing **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Configurations** > **All Configurations** > **Quota**.


**Step 6** Check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 7](#).

#### Collect fault information.

**Step 7** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

**Step 8** Select **ZooKeeper** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.84 ALM-13007 Available ZooKeeper Client Connections Are Insufficient

### Description

The system periodically detects the number of active processes between the ZooKeeper client and the ZooKeeper server every 60 seconds. This alarm is generated when the number of connections exceeds the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13007	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the host name for which the alarm is generated.
ClientIP	Specifies the client IP address.
ServerIP	Specifies the server IP address.
Trigger Condition	Specifies the cause of the alarm.

### Impact on the System


A large number of processes are connected to ZooKeeper, and the number of ZooKeeper connections is used up. As a result, services of upstream components (such as Yarn, Flink, and Spark) are abnormal.

### Possible Causes


A large number of client processes are connected to ZooKeeper. The thresholds are not appropriate.

## Procedure

**Check whether there are a large number of client processes connected to ZooKeeper.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the displayed interface, click the drop-down button of **Available ZooKeeper Client Connections Are Insufficient**. Confirm the node IP address of the host for which the alarm is generated in the Location Information.
- Step 2** Open the ZooKeeper service interface, click **Resource** to enter the **Resource** page, and check whether the number of connections of the client with the IP address specified by **Number of Connections (By Client IP Address)** is large.
- If it is, go to **Step 3**.
  - If it is not, go to **Step 4**.
- Step 3** Check whether connection leakage occurs on the client process.
- Step 4** Click  in the **Number of Connections (by Client IP Address)** to enter the **Thresholds** page, and click **Modify** under **Operation**. Increase the threshold by referring to the value of **maxClientCnxns** by choosing **Cluster > Name of the desired cluster > Services > ZooKeeper > Configurations > All Configurations > quorumpeer**.
- Step 5** Check whether the alarm is cleared.
- If it is, no further action is required.
  - If it is not, go to **Step 6**.

**Collect fault information.**

- Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 7** Select **ZooKeeper** in the required cluster from the **Service**.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.85 ALM-13008 ZooKeeper Znode Usage Exceeds the Threshold

### Description

The system checks the level-2 Znode status in the ZooKeeper data directory every hour. This alarm is generated when the system detects that the level-2 Znode usage exceeds the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13008	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
ServiceDirectory	Specifies the directory for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Trigger Condition	Specifies the cause of the alarm.

### Impact on the System

When a large amount of data is written to the ZooKeeper data directory space, ZooKeeper cannot provide services for external systems. As a result, services of upstream components (such as Yarn, Flink, and Spark) that depend on the alarm directory are abnormal.

### Possible Causes

- A large amount of data is written to the ZooKeeper data directory.
- The user-defined threshold is inappropriate.

### Procedure

**Check whether a large amount of data is written into the directory for which the alarm is generated.**


- Step 1** Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper**, and click **Resource**. Click **By Znode quantity in Used Resources (By Second-Level Znode)**, and check whether a large amount of data is written to the top Znode.
- If yes, go to **Step 2**.
  - If no, go to **Step 4**.
- Step 2** Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, select **Location** from the drop-down list box next to **ALM-13008 ZooKeeper Znode Quantity Usage Exceeds Threshold**, and obtain the Znode path in **ServiceDirectory**.
- Step 3** Log in to the ZooKeeper client as a cluster user and delete unnecessary data from the Znode corresponding to the alarm.
- Step 4** Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Configurations** > **All Configurations**, and search for **max.znode.count**, which is the maximum number of ZooKeeper directories. The alarm threshold is 80% of this parameter. Increase the value of this parameter, click **Save**, and restart the service for the configuration to take effect.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

- Step 5** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 6**.

**Collect fault information.**

- Step 6** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.
- Step 7** Select **ZooKeeper** in the required cluster from the **Service**.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.86 ALM-13009 ZooKeeper Znode Capacity Usage Exceeds the Threshold

### Alarm Description

The system checks the level-2 ZNode status in the ZooKeeper data directory every hour. This alarm is generated when the system detects that the capacity usage exceeds the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
13009	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
ServiceDirectory	Specifies the directory for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

ZooKeeper cannot provide services for external systems, and the services of upstream components (such as Yarn, Flink, and Spark) that depend on the alarm directory are abnormal.

### Possible Causes

- A large volume of data has been written to the ZooKeeper data directory.
- The threshold is improperly defined.

### Handling Procedure

**Check whether a large volume of data is written to the alarm directory.**

- Step 1** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. Click the drop-down list in the row containing **ALM-13009 ZooKeeper ZNode Capacity Usage Exceeds the Threshold**, and find the ZNode for which the alarm is generated in the **Location** area.
- Step 2** Choose **Cluster > Services > ZooKeeper**. On the page that is displayed, click the **Resource** tab. In the **Used Resources (By Second-Level ZNode)** area, click **By capacity** and check whether a large amount of data is written to the top-level ZNode directory.
- If yes, record the directory to which a large amount of data is written and go to **Step 3**.
  - If no, go to **Step 5**.
- Step 3** Check whether data in the directory can be deleted.

---

**NOTICE**

Deleting data from ZooKeeper is a high-risk operation. Exercise caution when performing this operation.

---


- If yes, go to **Step 4**.
  - If no, go to **Step 5**.
- Step 4** Log in to the ZooKeeper client and delete unnecessary data from the directory to which a large amount of data is written.
1. Log in to the ZooKeeper client installation directory, for example, **/opt/client**, and configure environment variables.  
**cd /opt/client**  
**source bigdata\_env**
  2. Run the following command to authenticate the user (skip this step for a cluster in normal mode):  
**kinit Component service user**
  3. Run the following command to log in to the client tool:  
**zkCli.sh -server <Service IP address of the node where any ZooKeeper instance resides>:<Client port>**
  4. Run the following command to delete unnecessary data:  
**delete Path of the file to be deleted**
- Step 5** Log in to FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > ZooKeeper > Configurations > All Configurations**, and search for **max.data.size**. The value of **max.data.size** is the maximum capacity quota of the ZooKeeper directory. The unit is byte. Search for the **GC\_OPTS** configuration item and check the value of **Xmx**.
- Step 6** Compare the values of **max.data.size** and **Xmx\*0.65**. The threshold is the smaller value multiplied by 80%. You can change the values of **max.data.size** and **Xmx\*0.65** to increase the threshold.
- Step 7** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect the fault information.**

**Step 8** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 9** Expand the **Service** drop-down list, and select **ZooKeeper** for the target cluster.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.87 ALM-13010 Znode Usage of a Directory with Quota Configured Exceeds the Threshold

### Description

The system checks the Znode usage of all service directories with quota configured every hour. This alarm is generated when the system detects that the level-2 Znode usage exceeds the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13010	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.



Name	Meaning
ServiceDirectory	Specifies the directory for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Trigger Condition	Specifies the cause of the alarm.

## Impact on the System

When a large amount of data is written to the ZooKeeper data directory space, ZooKeeper cannot provide services for external systems. As a result, services of upstream components (such as Yarn, Flink, and Spark) that depend on the alarm directory are abnormal.

## Possible Causes

- A large amount of data is written to the ZooKeeper data directory.
- The user-defined threshold is inappropriate.

## Procedure

**Check whether a large amount of data is written into the directory for which the alarm is generated.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Confirm the Znode for which the alarm is generated in **Location** of this alarm.
- Step 2** Choose **Cluster > Name of the desired cluster > Services > ZooKeeper** and click **Resource**. In **Used Resources (By Second-Level Znode)**, check whether a large amount of data is written into the top Znode.
  - If yes, go to **Step 3**.
  - If no, go to **Step 5**.
- Step 3** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, select Location from the drop-down list box next to **ALM-13010 Znode Usage of a Directory with Quota Configured Exceeds the Threshold**, and obtain the Znode path in ServiceDirectory.
- Step 4** Log in to the ZooKeeper client as a cluster user and delete unwanted data in the Znode for which the alarm is generated.
- Step 5** Log in to FusionInsight Manager, and choose **Cluster > Name of the desired cluster > Services > Component of the top Znode for which the alarm is generated**. Choose **Configurations > All Configurations**, search for **zk.quota.number**, increase its value, click **Save**.

**NOTICE**

If the Component of the top Znode for which the alarm is generated is ClickHouse, change the value of `clickhouse.zookeeper.quota.node.count`.


**Step 6** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **ZooKeeper** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.88 ALM-14000 HDFS Service Unavailable

### Description

The system checks the NameService service status every 60 seconds. This alarm is generated when all the NameService services are abnormal and the system considers that the HDFS service is unavailable.

This alarm is cleared when at least one NameService service is normal and the system considers that the HDFS service recovers.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14000	Critical	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

HDFS fails to provide services for HDFS service-based upper-layer components, such as HBase and MapReduce. As a result, users cannot read or write files.

## Possible Causes

- The ZooKeeper service is abnormal.
- All NameService services are abnormal.
- The number of service requests is too large, and the HDFS health check fails to read and write files.
- The health check fails due to HDFS FullGC.

## Procedure

### Check the ZooKeeper service status.

**Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the Alarm page, check whether **ALM-13000 ZooKeeper Service Unavailable** is reported.

- If yes, go to [Step 2](#).
- If no, go to [Step 4](#).

**Step 2** See **ALM-13000 ZooKeeper Service Unavailable** to rectify the health status of ZooKeeper fault and check whether the **Running Status** of the ZooKeeper service restores to **Normal**.

- If yes, go to [Step 3](#).
- If no, go to [Step 13](#).

**Step 3** On the **O&M > Alarm > Alarms** page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

### Handle the NameService service exception alarm.

**Step 4** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the Alarms page, check whether **ALM-14010 NameService Service Unavailable** is reported.

- If yes, go to [Step 5](#).
- If no, go to [Step 7](#).

**Step 5** See **ALM-14010 NameService Service Unavailable** to handle the abnormal NameService services and check whether each NameService service exception alarm is cleared.

- If yes, go to [Step 6](#).
- If no, go to [Step 13](#).

**Step 6** On the **O&M > Alarm > Alarms** page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Check whether the HDFS health check fails to read or write files due to a large number of service requests.**

**Step 7** On FusionInsight Manager, choose **O&M > Alarm > Alarms**, and check whether **ALM-14021 NameNode Average RPC Processing Time Exceeds the Threshold** or **ALM-14022 NameNode Average RPC Queuing Time Exceeds the Threshold** is generated.

- If yes, go to [Step 8](#).
- If no, go to [Step 10](#).

**Step 8** Rectify the abnormal NameServices by following the handling methods of **ALM-14021 NameNode Average RPC Processing Time Exceeds the Threshold** and **ALM-14022 NameNode Average RPC Queuing Time Exceeds the Threshold**. Then, check whether the alarms are cleared.

- If yes, go to [Step 9](#).
- If no, go to [Step 13](#).

**Step 9** On the **O&M > Alarm > Alarms** page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

**Check whether the health check fails due to HDFS FullGC.**

**Step 10** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the Alarms page, check whether **ALM-14014 NameNode GC Time Exceeds the Threshold** is reported.

- If yes, go to [Step 11](#).
- If no, go to [Step 13](#).

**Step 11** See **ALM-14014 NameNode GC Time Exceeds the Threshold** to handle the abnormal NameService services and check whether each NameService service exception alarm is cleared.

- If yes, go to [Step 12](#).
- If no, go to [Step 13](#).

**Step 12** On the **O&M > Alarm > Alarms** page, check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 13](#).

**Collect fault information.**

**Step 13** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 14** Select the following nodes in the required cluster from the **Service**:

- ZooKeeper
- HDFS

**Step 15** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 16** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.89 ALM-14001 HDFS Disk Usage Exceeds the Threshold

## Description

The system checks the HDFS disk usage every 30 seconds and compares the actual HDFS disk usage with the threshold. The HDFS disk usage indicator has a default threshold, this alarm is generated when the value of the disk usage of a Hadoop distributed file system (HDFS) indicator exceeds the threshold.

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS**.

When the **Trigger Count** is 1, this alarm is cleared when the value of the disk usage of HDFS cluster indicator is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the value of the disk usage of HDFS cluster indicator is less than or equal to 90% of the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14001	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

Writing Hadoop distributed file system (HDFS) data is affected.

## Possible Causes

The disk space configured for the HDFS cluster is insufficient.

## Procedure

**Check the disk capacity and delete unnecessary files.**

- Step 1** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**.
- Step 2** Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Disk**, and select **Percentage of HDFS Capacity** to check whether the HDFS disk usage exceeds the threshold (80% by default).
  - If yes, go to **Step 3**.
  - If no, go to **Step 11**.
- Step 3** In the **Basic Information** area, click the **NameNode(Active)** of the failure NameService and the HDFS WebUI page is displayed.


### NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

**Step 4** On the HDFS web user interface (WebUI), click **Datanodes** tab. In the **Block pool used** column, view the disk usage of all DataNodes to check whether the disk usage of any DataNode exceeds the threshold.

- If yes, go to [Step 6](#).
- If no, go to [Step 11](#).

**Figure 7-69** Datanode Information



**Step 5** Log in to the MRS client node as user **root**.

**Step 6** Run **cd /opt/client** to switch to the client installation directory, and run **source bigdata\_env**. If the cluster uses the security mode, perform security authentication. Run **kinit hdfs** and enter the password as prompted. Please obtain the password from the administrator.

**Step 7** Run the **hdfs dfs -rm -r file or directory** command to delete unnecessary files.

**Step 8** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Expand the system.**

**Step 9** Expand the disk capacity.

**Step 10** Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 11](#).

**Collect fault information.**

**Step 11** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 12** Select the following nodes in the required cluster from the **Service**:

- ZooKeeper
- HDFS

**Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 14** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.90 ALM-14002 DataNode Disk Usage Exceeds the Threshold

### Alarm Description

The system checks the DataNode disk usage every 30 seconds and compares the actual disk usage with the threshold. A default threshold range is provided for the DataNode disk usage. This alarm is generated when the DataNode disk usage exceeds the threshold.

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS**.

If **Trigger Count** is **1**, this alarm is cleared when the DataNode disk usage is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the DataNode disk usage is less than or equal to 80% of the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14002	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.



Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Insufficient disk space will impact data write to HDFS.

## Possible Causes

- The disk space configured for the HDFS cluster is insufficient.
- Data skew occurs among DataNodes.

## Handling Procedure

**Check whether the cluster disk capacity is insufficient.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**, and check whether the **ALM-14001 HDFS Disk Usage Exceeds the Threshold** alarm exists.
- If yes, go to [Step 2](#).
  - If no, go to [Step 4](#).
- Step 2** Handle the alarm by following the instructions in **ALM-14001 HDFS Disk Usage Exceeds the Threshold** and check whether the alarm is cleared.
- If yes, go to [Step 3](#).
  - If no, go to [Step 11](#).
- Step 3** Choose **O&M > Alarm > Alarms** and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 4](#).

**Check the balance status of DataNodes.**

- Step 4** On FusionInsight Manager, choose **Hosts**. Check whether the number of DataNodes on each rack is almost the same. If the difference is large, adjust the racks to which DataNodes belong to ensure that the number of DataNodes on each rack is almost the same. Restart the HDFS service for the settings to take effect.

### NOTE

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

**Step 5** Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**.

**Step 6** In the **Basic Information** area, click **NameNode(Active)**. The HDFS web UI is displayed.

 **NOTE**

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

**Step 7** In the **Summary** area of the HDFS web UI, check whether the value of **Max** is 10% greater than that of **Median** in **DataNodes usages**.

- If yes, go to [Step 8](#).
- If no, go to [Step 11](#).

**Step 8** Balance skewed data in the cluster. Log in to the MRS client as user **root**. If the cluster is in normal mode, run the **su - omm** command to switch to user **omm**. Run the **cd** command to go to the client installation directory and run the **source bigdata\_env** command. If the cluster uses the security mode, perform security authentication. Run **kinit hdfs** and enter the password as prompted. Obtain the password from the MRS cluster administrator.

**Step 9** Run the following command to balance data distribution:

```
hdfs balancer -threshold 10
```


**Step 10** Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Collect the fault information.**

**Step 11** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 12** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

**Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 14** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.91 ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold

### Alarm Description

The system checks the lost blocks every 30 seconds and compares the actual lost blocks with the threshold. The lost blocks indicator has a default threshold. This alarm is generated when the number of lost HDFS blocks exceeds the threshold.

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS**.

If **Trigger Count** is **1**, this alarm is cleared when the value of lost HDFS blocks is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the value of lost HDFS blocks is less than or equal to 90% of the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14003	Major <b>NOTE</b> The alarm severity in MRS 3.1.5 is <b>Critical</b> .	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Data stored in HDFS is lost. HDFS may enter the security mode and cannot provide write services. Lost block data cannot be restored.

## Possible Causes

- The DataNode instance is abnormal.
- Data is deleted.

## Handling Procedure

### Check the DataNode instance.

**Step 1** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance**.

**Step 2** Check whether the **Running Status** of all DataNode instance is **Normal**.

- If yes, go to [Step 11](#).
- If no, go to [Step 3](#).

**Step 3** Restart the DataNode instance and check whether the DataNode instance restarts successfully.

#### NOTE

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.

- If yes, go to [Step 4](#).
- If no, go to [Step 5](#).

**Step 4** Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

### Delete the damaged file.

**Step 5** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **NameNode(Active)**. On the WebUI page of the HDFS, view the information about lost blocks.

#### NOTE

- If a block is lost, a line in red is displayed on the WebUI.
- By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

**Step 6** The user checks whether the file containing the lost data block is useful.

#### NOTE

Files generated in directories **/mr-history**, **/tmp/hadoop-yarn**, and **/tmp/logs** during MapReduce task execution are unnecessary.

- If yes, go to [Step 7](#).

- If no, go to [Step 8](#).

**Step 7** The user checks whether the file containing the lost data block is backed up.

- If yes, go to [Step 8](#).
- If no, go to [Step 11](#).

**Step 8** Log in to the HDFS client as user **root**. The user password is defined by the user before the installation. Contact the MRS cluster administrator to obtain the password. Run the following commands:

- Security mode:  
`cd Client installation directory`  
`source bigdata_env`  
`kinit hdfs`
- Normal mode:  
`su - omm`  
`cd Client installation directory`  
`source bigdata_env`

**Step 9** On the node client, run **hdfs fsck / -delete** to delete the lost file. If the file where the lost block is located is a useful file, you need to write the file again to restore the data.

 **NOTE**

Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.


**Step 10** Choose **O&M > Alarm > Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Collect the fault information.**

**Step 11** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 12** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

**Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 14** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.92 ALM-14006 Number of HDFS Files Exceeds the Threshold

### Alarm Description

The system periodically checks the number of HDFS files every 30 seconds and compares the number of HDFS files with the threshold. This alarm is generated when the system detects that the number of HDFS files exceeds the threshold.

If **Trigger Count** is **1**, this alarm is cleared when the number of HDFS files is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the number of HDFS files is less than or equal to 90% of the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14006	Minor	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

If there are too many HDFS files, the HDFS system may respond slowly or the disk space may be used up.

### Possible Causes

The number of HDFS files exceeds the threshold.

## Handling Procedure

### Check the number of files in the system.

- Step 1** On FusionInsight Manager, check the number of HDFS files. Specifically, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **File and Block**, and select **HDFS File** and **Total Blocks**.
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** > **All Configurations**, and search for the **GC\_OPTS** parameter under **NameNode**.
- Step 3** Configure the threshold of the number of configuration file objects. Specifically, change the value of **Xmx** (GB) in the **GC\_OPTS** parameter. The threshold (specified by *y*) is calculated as follows:  $y = 0.2007 \times Xmx - 0.6312$ , where *x* indicates the memory capacity *Xmx* (GB) and *y* indicates the number of files (unit: kW). Adjust the memory size as required.
- Step 4** Confirm that the value of **GC\_PROFILE** is **custom** so that the **GC\_OPTS** configuration takes effect. Click **Save** and choose **More** > **Restart Instance** to restart the service.

#### NOTE

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

- Step 5** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).

### Check whether needless files exist in the system.

- Step 6** Log in to the HDFS client as user **root**. Run **cd** to switch to the client installation directory, and run **source bigdata\_env** to configure the environment variables.

If the cluster uses the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the MRS cluster administrator.

- Step 7** Run **hdfs dfs -ls file or directory** to check whether the files in the directory can be deleted.
- If yes, go to [Step 8](#).
  - If no, go to [Step 9](#).

- Step 8** Run the **hdfs dfs -rm -r file or directory path** command. After deleting unnecessary files, wait until the files are retained in the recycle bin for a period longer than the value of **fs.trash.interval** on the NameNode. Then check whether the alarm is cleared.

#### NOTE

Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.


- If yes, no further action is required.

- If no, go to [Step 9](#).

**Collect the fault information.**

**Step 9** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 10** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

### Configuration rules of the NameNode JVM parameter

Default value of the NameNode JVM parameter **GC\_OPTS**:

```
-Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M -
XX:MetaspaceSize=128M -XX:MaxMetaspaceSize=128M -
XX:+UseConcMarkSweepGC -XX:+CMSParallelRemarkEnabled -
XX:CMSInitiatingOccupancyFraction=65 -XX:+PrintGCDetails -
Dsun.rmi.dgc.client.gcInterval=0x7FFFFFFF -
Dsun.rmi.dgc.server.gcInterval=0x7FFFFFFF -XX:-
OmitStackTraceInFastThrow -XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation
-XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=1M -
Djdk.tls.ephemeralDHKeySize=3072 -
Djdk.tls.rejectClientInitiatedRenegotiation=true -Djava.io.tmpdir=$
{Bigdata_tmp_dir}
```

The number of NameNode files is proportional to the used memory size of the NameNode. When file objects change, you need to change **-Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M** in the default value. The following table lists the reference values.

**Table 7-101** NameNode JVM configuration

Number of File Objects	Reference Value
10,000,000	-Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
20,000,000	-Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G



Number of File Objects	Reference Value
50,000,000	-Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G
100,000,000	-Xms64G -Xmx64G -XX:NewSize=6G -XX:MaxNewSize=6G
200,000,000	-Xms96G -Xmx96G -XX:NewSize=9G -XX:MaxNewSize=9G
300,000,000	-Xms164G -Xmx164G -XX:NewSize=12G -XX:MaxNewSize=12G

## 7.12.93 ALM-14007 NameNode Heap Memory Usage Exceeds the Threshold

### Description

The system checks the HDFS NameNode Heap Memory usage every 30 seconds and compares the actual Heap memory usage with the threshold. The HDFS NameNode Heap Memory usage has a default threshold. This alarm is generated when the HDFS NameNode Heap Memory usage exceeds the threshold.

You can change the threshold in **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS**.

When the **Trigger Count** is 1, this alarm is cleared when the HDFS NameNode Heap memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the HDFS NameNode Heap memory usage is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14007	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

The HDFS NameNode Heap Memory usage is too high, which affects the data read/write performance of the HDFS.

## Possible Causes

The HDFS NameNode Heap Memory is insufficient.

## Procedure

### Delete unnecessary files.

**Step 1** Log in to the HDFS client as user **root**. Run **cd** to switch to the client installation directory, and run **source bigdata\_env**.

If the cluster uses the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

**Step 2** Run the **hdfs dfs -rm -r file or directory** command to delete unnecessary files.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

### Check the NameNode JVM memory usage and configuration.

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS**.

**Step 5** In the **Basic Information** area, click **NameNode(Active)** to go to the HDFS WebUI.

### NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

- Step 6** On the HDFS WebUI, click the **Overview** tab. In **Summary**, check the numbers of files, directories, and blocks in the HDFS.
- Step 7** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations > All Configurations**. In **Search**, enter **GC\_OPTS** to check the **GC\_OPTS** memory parameter of **HDFS->NameNode**.

**Adjust the configuration in the system.**

- Step 8** Check whether the memory is configured properly based on the number of files in [Step 6](#) and the NameNode Heap Memory parameters in [Step 7](#).
- If yes, go to [Step 9](#).
  - If no, go to [Step 11](#).

 **NOTE**

The recommended mapping between the number of HDFS file objects (filesystem objects = files + blocks) and the JVM parameters configured for NameNode is as follows:

- If the number of file objects reaches 10,000,000, you are advised to set the JVM parameters as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the number of file objects reaches 20,000,000, you are advised to set the JVM parameters as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G
- If the number of file objects reaches 50,000,000, you are advised to set the JVM parameters as follows: -Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G
- If the number of file objects reaches 100,000,000, you are advised to set the JVM parameters as follows: -Xms64G -Xmx64G -XX:NewSize=6G -XX:MaxNewSize=6G
- If the number of file objects reaches 200,000,000, you are advised to set the JVM parameters as follows: -Xms96G -Xmx96G -XX:NewSize=9G -XX:MaxNewSize=9G
- If the number of file objects reaches 300,000,000, you are advised to set the JVM parameters as follows: -Xms164G -Xmx164G -XX:NewSize=12G -XX:MaxNewSize=12G

- Step 9** Modify the heap memory parameters of the NameNode based on the mapping between the number of file objects and the memory. Click **Save** and choose **Dashboard > More > Restart Service**.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.


- Step 10** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 11](#).

**Collect fault information.**

- Step 11** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

- Step 12** Select the following nodes in the required cluster from the **Service**:

- ZooKeeper
- HDFS

- Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 14** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.94 ALM-14008 DataNode Heap Memory Usage Exceeds the Threshold

## Description

The system checks the HDFS DataNode Heap Memory usage every 30 seconds and compares the actual Heap Memory usage with the threshold. The HDFS DataNode Heap Memory usage has a default threshold. This alarm is generated when the HDFS DataNode Heap Memory usage exceeds the threshold.

You can change the threshold in **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS**.

When the **Trigger Count** is 1, this alarm is cleared when the HDFS DataNode Heap Memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the HDFS DataNode Heap Memory usage is less than or equal to 90% of the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14008	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

The HDFS DataNode Heap Memory usage is too high, which affects the data read/write performance of the HDFS.

## Possible Causes

The HDFS DataNode Heap Memory is insufficient.

## Procedure

### Delete unnecessary files.

**Step 1** Log in to the HDFS client as user **root**. Run **cd** to switch to the client installation directory, and run **source bigdata\_env**.

If the cluster uses the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

**Step 2** Run the **hdfs dfs -rm -r file or directory** command to delete unnecessary files.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

### Check the DataNode JVM memory usage and configuration.

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS**.

**Step 5** In the **Basic Information** area, click **NameNode(Active)** to go to the HDFS WebUI.

### NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

**Step 6** On the HDFS WebUI, click the **DataNodes** tab, and check the number of blocks of all DataNodes related to the alarm.

**Step 7** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations > All Configurations**. In **Search**, enter **GC\_OPTS** to check the GC\_OPTS memory parameter of **HDFS->DataNode**.

### Adjust the configuration in the system.

**Step 8** Check whether the memory is configured properly based on the number of block in [Step 6](#) and the DataNode Heap Memory parameters in [Step 7](#).

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

#### NOTE

The mapping between the average number of blocks of a DataNode instance and the DataNode memory is as follows:

- If the average number of blocks of a DataNode instance reaches 2,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the average number of blocks of a DataNode instance reaches 5,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

**Step 9** Modify the heap memory parameters of the DataNode based on the mapping between the number of blocks and the memory. Click **Save** and choose **Dashboard > More > Restart Service**.

#### NOTE

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.


**Step 10** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

### Collect fault information.

**Step 11** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 12** Select **HDFS** in the required cluster from the **Service**.

**Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 14** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.95 ALM-14009 Number of Dead DataNodes Exceeds the Threshold

### Description

The system periodically detects the number of dead DataNodes in the HDFS cluster every 30 seconds, and compares the number with the threshold. The number of DataNodes in the Dead state has a default threshold. This alarm is generated when the number exceeds the threshold.

You can change the threshold in **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS**.

When the **Trigger Count** is 1, this alarm is cleared when the number of Dead DataNodes is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the number of Dead DataNodes is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14009	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

DataNodes that are in the Dead state cannot provide HDFS services. As a result, users cannot read or write files.

## Possible Causes

- DataNodes are faulty or overloaded.
- The network between the NameNode and the DataNode is disconnected or busy.
- NameNodes are overloaded.
- The NameNodes are not restarted after the DataNode is deleted.

## Procedure

### Check whether DataNodes are faulty.

**Step 1** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**. The **HDFS Status** page is displayed.

**Step 2** In the **Basic Information** area, click **NameNode(Active)** to go to the HDFS WebUI.

#### NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

**Step 3** On the HDFS WebUI, click the **Datanodes** tab. In the **In operation** area, click **Filter** to check whether **down** is in the drop-down list.

- If yes, select **down**, record the information about the filtered DataNodes, and go to **Step 4**.
- If no, go to **Step 8**.

**Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance** to check whether recorded DataNodes exist in the instance list.

- If all recorded DataNodes exist, go to **Step 5**.
- If none of the recorded DataNodes exists, go to **Step 6**.
- If some of the recorded DataNodes exist, go to **Step 7**.

**Step 5** Locate the DataNode instance, click **More** > **Restart Instance** to restart it and check whether the alarm is cleared.

#### NOTE

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.

- If yes, no further action is required.
- If no, go to **Step 8**.

**Step 6** Select all NameNode instances, choose **More** > **Instance Rolling Restart** to restart them and check whether the alarm is cleared.



- If yes, no further action is required.
- If no, go to [Step 16](#).

**Step 7** Select all NameNode instances, choose **More > Instance Rolling Restart** to restart them. Locate the DataNode instance, click **More > Restart Instance** to restart it and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Check the status of the network between the NameNode and the DataNode.**

**Step 8** Log in to the faulty DataNode on the management page as user **root**, and run the **ping IP address of the NameNode** command to check whether the network between the DataNode and the NameNode is abnormal.

On the FusionInsight Manager page, choose **Cluster > Name of the desired cluster > Services > HDFS > Instance**. In the instance list, view the service plane IP address of the faulty DataNode.

- If yes, go to [Step 9](#).
- If no, go to [Step 10](#).

**Step 9** Rectify the network fault, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

**Check whether the DataNode is overloaded.**

**Step 10** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and check whether the alarm **ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold** exists.

- If yes, go to [Step 11](#).
- If no, go to [Step 13](#).

**Step 11** See **ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold** to handle the alarm and check whether the alarm is cleared.

- If yes, go to [Step 12](#).
- If no, go to [Step 13](#).

**Step 12** Check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to [Step 13](#).

**Check whether the NameNode is overloaded.**

**Step 13** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and check whether the alarm **ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold** exists.

- If yes, go to [Step 14](#).
- If no, go to [Step 16](#).

**Step 14** See **ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold** to handle the alarm and check whether the alarm is cleared.

- If yes, go to [Step 15](#).
- If no, go to [Step 16](#).


**Step 15** Check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to [Step 16](#).

**Collect fault information.**

**Step 16** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 17** Select **HDFS** in the required cluster from the **Service**.

**Step 18** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 19** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.96 ALM-14010 NameService Service Is Abnormal

### Alarm Description

The system checks the NameService service status every 180 seconds. This alarm is generated when the NameService service is unavailable.

This alarm is cleared when the NameService service recovers.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14010	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.

## Impact on the System

HDFS fails to provide services for upper-layer components based on the NameService service, such as HBase and MapReduce. As a result, users cannot read or write files.

## Possible Causes

- The KrbServer service is abnormal.
- The JournalNode is faulty.
- The DataNode is faulty.
- The disk capacity is insufficient.
- The NameNode enters safe mode.

## Handling Procedure

### Check the KrbServer service status.

- Step 1** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**.
- Step 2** Check whether the KrbServer service exists.
- If yes, go to [Step 3](#).
  - If no, go to [Step 6](#).
- Step 3** Click **KrbServer**.
- Step 4** Click **Instances**. On the KrbServer management page, select the faulty instance, and choose **More** > **Restart Instance**. Check whether the instance successfully restarts.
- If yes, go to [Step 5](#).
  - If no, go to [Step 24](#).
- Step 5** Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).

### Check the JournalNode instance status.

**Step 6** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**.

**Step 7** Choose **HDFS** > **Instances**.

**Step 8** Check whether the **Running Status** of the JournalNode is **Normal**.

- If yes, go to [Step 11](#).
- If no, go to [Step 9](#).

**Step 9** Select the faulty JournalNode, and choose **More** > **Restart Instance**. Check whether the JournalNode successfully restarts.

 **NOTE**

If the number of JournalNode instances restarted exceeds one third of the total number of JournalNodes, the HDFS service may be faulty.

- If yes, go to [Step 10](#).
- If no, go to [Step 24](#).

**Step 10** Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Check the DataNode instance status.**

**Step 11** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**.

**Step 12** Click **Instances** and check whether **Running Status** of all DataNodes is **Normal**.

- If yes, go to [Step 15](#).
- If no, go to [Step 13](#).

**Step 13** Click **Instances**. On the DataNode management page, select the faulty instance, and choose **More** > **Restart Instance**. Check whether the DataNode successfully restarts.

 **NOTE**

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.

- If yes, go to [Step 14](#).
- If no, go to [Step 15](#).

**Step 14** Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 15](#).

**Check disk status.**

**Step 15** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Host**.

**Step 16** In the **Disk** column, check whether the disk space is insufficient.

- If yes, go to [Step 17](#).
- If no, go to [Step 19](#).

**Step 17** Expand the disk capacity.

**Step 18** Choose **O&M > Alarm > Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 19](#).

**Check whether NameNode is in the safe mode.**

**Step 19** On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > HDFS**. Click **NameNode(Active)** of the abnormal NameService. The NameNode web UI is displayed.

#### NOTE


By default, the admin user does not have the management rights of other components. If the page cannot be opened or the content is not completely displayed due to insufficient permission when you access the native page of a component, you can manually create a user with the management rights of the corresponding component to log in to the component.

**Step 20** On the NameNode web UI, check whether "Safe mode is ON." is displayed.

Information behind **Safe mode is ON** is alarm information and is displayed based actual conditions.

- If yes, go to [Step 21](#).
- If no, go to [Step 24](#).

**Figure 7-70** Overview



Namespace:	hacluster
Namenode ID:	3
Started:	Tue Sep 28 15:37:40 +0800 2021
Version:	3
Compiled:	Wed Jun 23 11:06:00 +0800 2021 by root from (HEAD detached at 1df7014cdd)
Cluster ID:	myhacluster
Block Pool ID:	BP-1081137484-192.168.0.146-1632814642995

### Summary

Security is off.  
Safemode is off.  
922 files and directories, 769 blocks (769 replicated blocks, 0 erasure coded block groups) = 1,691 total filesystem object(s).  
Heap Memory used 329.67 MB of 3.95 GB Heap Memory. Max Heap Memory is 7.82 GB.  
Non Heap Memory used 106.84 MB of 108.88 MB Committed Non Heap Memory. Max Non Heap Memory is 488 MB.

**Step 21** Log in to the client as user **root**. Run the **cd** command to go to the client installation directory and run the **source bigdata\_env** command. If the cluster uses the security mode, perform security authentication. Run the **kinit hdfs** command and enter the password as prompted. The password can be obtained from the MRS cluster administrator. If the cluster uses the non-security mode, log in as user **omm** and run the command. Ensure that user **omm** has the client execution permission.

**Step 22** Run `hdfs dfsadmin -safemode leave`.

**Step 23** Choose **O&M > Alarm > Alarms** and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 24](#).

**Collect the fault information.**

**Step 24** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 25** In the **Service** area, select the following nodes of the desired cluster.

- ZooKeeper
- HDFS

**Step 26** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 27** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.97 ALM-14011 DataNode Data Directory Is Not Configured Properly

## Description

The DataNode parameter `dfs.datanode.data.dir` specifies DataNode data directories. This alarm is generated when a configured data directory cannot be created, a data directory uses the same disk as other critical directories in the system, or multiple directories use the same disk immediately.

This alarm is cleared when the DataNode data directory is configured properly and this DataNode for which the alarm is generated is restarted.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14011	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

If the DataNode data directory is mounted to the root directory or a critical directory, the disk space of the root directory or critical directory will be used up after long time running and the system will be faulty.

If the DataNode data directory is not configured properly, HDFS performance will deteriorate.

## Possible Causes

- The DataNode data directory fails to be created.
- The DataNode data directory uses the same disk with critical directories, such as / or /boot.
- Multiple directories in the DataNode data directory use the same disk.

## Procedure

**Check the alarm cause and information about the DataNode for which the alarm is generated.**

**Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.

**Step 2** In **HostName** of **Location**, obtain the host name of the DataNode for which the alarm is generated.

**Delete directories that do not comply with the disk plan from the DataNode data directory.**

**Step 3** Choose **Cluster > Name of the desired cluster > Services > HDFS > Instance**. In the instance list, click the DataNode instance on the node for which the alarm is generated.

**Step 4** Click **Instance Configurations** and view the value of the DataNode parameter **dfs.datanode.data.dir**.

**Step 5** Check whether all DataNode data directories are consistent with the disk plan.

- If yes, go to [Step 6](#).
- If no, go to [Step 9](#).

**Step 6** Modify the DataNode parameter **dfs.datanode.data.dir** and delete the incorrect directories.

**Step 7** Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance** and restart the DataNode instance.

 **NOTE**

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.

**Step 8** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Step 9** Log in to the DataNode for which the alarm is generated as user **root**.

- If the alarm cause is "The DataNode data directory fails to be created", go to [Step 10](#).
- If the alarm cause is "The DataNode data directory uses the same disk with critical directories, such / or /boot", go to [Step 17](#).
- If the alarm cause is "Multiple directories in the DataNode data directory uses the same disk", go to [Step 21](#).

**Check whether the DataNode data directory fails to be created.**

**Step 10** Run the **su - omm** command to switch to user **omm**.

**Step 11** Run the **ls** command to check whether the directories exist in the DataNode data directory.

- If yes, go to [Step 26](#).
- If no, go to [Step 12](#).

**Step 12** Run the **mkdir data directory** command to create the directory and check whether the directory can be successfully created.

- If yes, go to [Step 24](#).
- If no, go to [Step 13](#).

**Step 13** On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** to check whether alarm **ALM-12017 Insufficient Disk Capacity** exists.

- If yes, go to [Step 14](#).
- If no, go to [Step 15](#).

**Step 14** Adjust the disk capacity and check whether alarm **ALM-12017 Insufficient Disk Capacity** is cleared. For details, see **ALM-12017 Insufficient Disk Capacity**.

- If yes, go to [Step 12](#).
- If no, go to [Step 15](#).

**Step 15** Check whether user **omm** has the **rwX** or **X** permission of all the upper-layer directories of the directory. (For example, for **/tmp/abc/**, user **omm** has the **X** permission for directory **tmp** and the **rwX** permission for directory **abc**.)



- If yes, go to [Step 24](#).
- If no, go to [Step 16](#).

**Step 16** Run the `chmod u+rwX path` or `chmod u+x path` command as user **root** to assign the **rwX** or **x** permission of these directories to user **omm**. Then go to [Step 12](#).

**Check whether the DataNode data directory use the same disk as other critical directories in the system.**

**Step 17** Run the `df` command to obtain the disk mounting information of each directory in the DataNode data directory.

**Step 18** Check whether the directories mounted to the disk are critical directories, such as `/` or `/boot`.

- If yes, go to [Step 19](#).
- If no, go to [Step 24](#).

**Step 19** Change the value of the DataNode parameter `dfs.datanode.data.dir` and delete the directories that use the same disk as critical directories.

**Step 20** Go to [Step 24](#).

**Check whether multiple directories in the DataNode data directory use the same disk.**

**Step 21** Run the `df` command to obtain the disk mounting information of each directory in the DataNode data directory. Record the mounted directory in the command output.

**Step 22** Modify the DataNode node parameters `dfs.datanode.data.dir` to reserve only one directory among the directories that mounted to the same disk directory.

**Step 23** Go to [Step 24](#).

**Restart the DataNode and check whether the alarm is cleared.**

**Step 24** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance** and restart the DataNode instance.

 **NOTE**

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.


**Step 25** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 26](#).

**Collect fault information.**

**Step 26** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

**Step 27** Select **HDFS** in the required cluster from the **Service**.

**Step 28** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 29** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.98 ALM-14012 JournalNode Is Out of Synchronization

## Description

On the active NameNode, the system checks the data consistency of all JournalNodes in the cluster every 5 minutes. This alarm is generated when the data on a JournalNode is inconsistent with the data on the other JournalNodes.

This alarm is cleared in 5 minutes after the data on JournalNodes is consistent.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14012	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService for which the alarm is generated.

## Impact on the System

When a JournalNode is working incorrectly, the data on the node becomes inconsistent with that on the other JournalNodes. If data on more than half of

JournalNodes is inconsistent, the NameNode cannot work correctly, making the HDFS service unavailable.

## Possible Causes

- The JournalNode instance does not exist (deleted or migrated).
- The JournalNode instance has not been started or has been stopped.
- The JournalNode instance is working incorrectly.
- The network of the JournalNode is unreachable.

## Procedure

### Check whether the JournalNode instance has been started up.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.
- Step 2** Check **Location** and obtain the IP address of the JournalNode for which the alarm is generated.
- Step 3** Choose **Cluster > Name of the desired cluster > Services > HDFS > Instance**. In the instance list, check whether the JournalNode instance exists on the node for which the alarm is generated.
- If yes, go to [Step 5](#).
  - If no, go to [Step 4](#).
- Step 4** Choose **O&M > Alarm > Alarms**. In the alarm list, click **Clear** in the **Operation** column of the alarm. In the dialog box that is displayed, click **OK**. No further action is needed.
- Step 5** Click the JournalNode instance and check whether its **Configuration Status** is **Synchronized**.
- If yes, go to [Step 8](#).
  - If no, go to [Step 6](#).
- Step 6** Select the JournalNode instance and choose **Start Instance** to start the instance.
- Step 7** After 5 minutes, check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 15](#).

### Check whether the JournalNode instance is working correctly.

- Step 8** Check whether **Running Status** of the JournalNode instance is **Normal**.
- If yes, go to [Step 11](#).
  - If no, go to [Step 9](#).
- Step 9** Select the JournalNode instance and choose **More > Restart Instance** to start the instance.

#### NOTE

If the number of JournalNode instances restarted exceeds one third of the total number of JournalNodes, the HDFS service may be faulty.

**Step 10** After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 15](#).

**Check whether the network of the JournalNode is reachable.**

**Step 11** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance** to check the service IP address of the active NameNode.

**Step 12** Log in to the active NameNode as user **root**.

**Step 13** Run the **ping** command to check whether a timeout occurs or the network is unreachable between the active NameNode and the JournalNode.

**ping** *service IP address of the JournalNode*

- If yes, go to [Step 14](#).
- If no, go to [Step 15](#).


**Step 14** Contact the network administrator to rectify the network fault and check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to [Step 15](#).

**Collect fault information.**

**Step 15** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

**Step 16** Select **HDFS** in the required cluster from the **Service**.

**Step 17** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 18** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.99 ALM-14013 Failed to Update the NameNode FsImage File

### Description

HDFS metadata is stored in the FsImage file of the NameNode data directory, which is specified by the **dfs.namenode.name.dir** configuration item. The standby NameNode periodically combines existing FsImage files and Editlog files stored in

the JournalNode to generate a new FsImage file, and then pushes the new FsImage file to the data directory of the active NameNode. This period is specified by the **dfs.namenode.checkpoint.period** configuration item of HDFS. The default value is 3600s, namely, one hour. If the FsImage file in the data directory of the active NameNode is not updated, the HDFS metadata combination function is abnormal and requires rectification.

On the active NameNode, the system checks the FsImage file information every five minutes. This alarm is generated when no FsImage file is generated within three combination periods.

This alarm is cleared when a new FsImage file is generated and pushed to the active NameNode, which indicates that the HDFS metadata combination function can be properly used.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14013	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService for which the alarm is generated.

## Impact on the System

If the FsImage file in the data directory of the active NameNode is not updated, the HDFS metadata combination function is abnormal and requires rectification. If it is not rectified, the Editlog files increase continuously after HDFS runs for a period. In this case, HDFS restart is time-consuming because a large number of Editlog files need to be loaded. In addition, this alarm also indicates that the standby NameNode is abnormal and the NameNode high availability (HA) mechanism becomes invalid. When the active NameNode is faulty, the HDFS service becomes unavailable.

## Possible Causes

- The standby NameNode is stopped.
- The standby NameNode instance is working incorrectly.
- The standby NameNode fails to generate a new FsImage file.
- Space of the data directory on the standby NameNode is insufficient.
- The standby NameNode fails to push the FsImage file to the active NameNode.
- Space of the data directory on the active NameNode is insufficient.

## Procedure

### Check whether the standby NameNode is stopped.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.
- Step 2** View **Location** and obtain the host name of the active NameNode for which the alarm is generated and name of the NameService where the active NameNode resides.
- Step 3** Choose **Cluster > Name of the desired cluster > Services > HDFS > Instance**, find the standby NameNode instance of the NameService in the instance list, and check whether its **Configuration Status** is **Synchronized**.
- If yes, go to **Step 6**.
  - If no, go to **Step 4**.
- Step 4** Select the standby NameNode instance, choose **Start Instance**, and wait until the startup is complete.
- Step 5** Wait for a NameNode metadata combination period and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 6**.

### Check whether the NameNode instance is working correctly.

- Step 6** Check whether **Running Status** of the standby NameNode instance is **Normal**.
- If yes, go to **Step 9**.
  - If no, go to **Step 7**.
- Step 7** Select the standby NameNode instance, choose **More > Restart Instance**, and wait until the startup is complete.

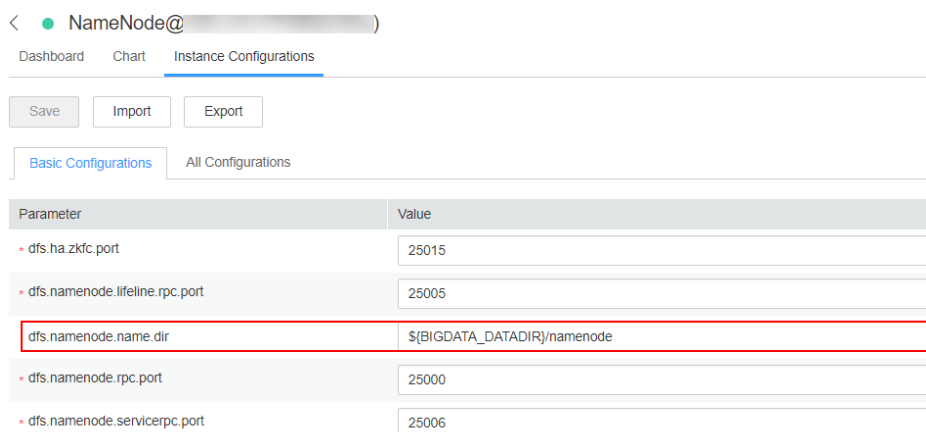
#### NOTE

Services are not affected after the standby NameNode is restarted.

- Step 8** Wait for a NameNode metadata combination period and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 30**.

### Check whether the standby NameNode fails to generate a new FsImage file.

- Step 9** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations > All Configurations**, and search and obtain the value of **dfs.namenode.checkpoint.period**. This value is the period of NameNode metadata combination.
- Step 10** Choose **Cluster > Name of the desired cluster > Services > HDFS > Instance** and obtain the service IP addresses of the active and standby NameNodes of the NameService for which the alarm is generated.
- Step 11** Click the **NameNode(XX,Standby)** and **Instance Configurations** to obtain the value of **dfs.namenode.name.dir**. This value is the FsImage storage directory of the standby NameNode.



- Step 12** Log in to the standby NameNode as user **root** or **omm**.
- Step 13** Go to the FsImage storage directory and check the generation time of the newest FsImage file.

```
cd Storage directory of the standby NameNode/current
stat -c %y $(ls -t | grep "fsimage_[0-9]*$" | head -1)
```

- Step 14** Run the **date** command to obtain the current system time.
- Step 15** Calculate the time difference between the generation time of the newest FsImage file and the current system time and check whether the time difference is greater than three times of the metadata combination period.
- If yes, go to **Step 16**.
  - If no, go to **Step 20**.

- Step 16** The metadata combination function of the standby NameNode is faulty. Run the following command to check whether the fault is caused by insufficient storage space.

Go to the FsImage storage directory and check the size of the newest FsImage file (in MB).

```
cd Storage directory of the standby NameNode/current
du -m $(ls -t | grep "fsimage_[0-9]*$" | head -1) | awk '{print $1}'
```

- Step 17** Run the following command to check the available disk space of the standby NameNode (in MB).

```
df -m ./ | awk 'END{print $4}'
```

**Step 18** Compare the Fslmage file size and the available disk space and determine whether another Fslmage file can be stored on the disk.

- If yes, go to [Step 7](#).
- If no, go to [Step 19](#).

**Step 19** Clear the redundant files on the disk where the directory resides to reserve sufficient space for metadata. After the clearance, wait for a NameNode metadata combination period and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 20](#).

**Check whether the standby NameNode fails to push the Fslmage file to the active NameNode.**

**Step 20** Log in to the standby NameNode as user **root**.

**Step 21** Run the **su - omm** command to switch to user **omm**.

**Step 22** Run the following command to check whether the standby NameNode can push the file to the active NameNode.

```
tmpFile=/tmp/tmp_test_$(date +%s)
```

```
echo "test" > $tmpFile
```

```
scp $tmpFile Service IP address of the active NameNode:/tmp
```

- If yes, go to [Step 24](#).
- If no, go to [Step 23](#).

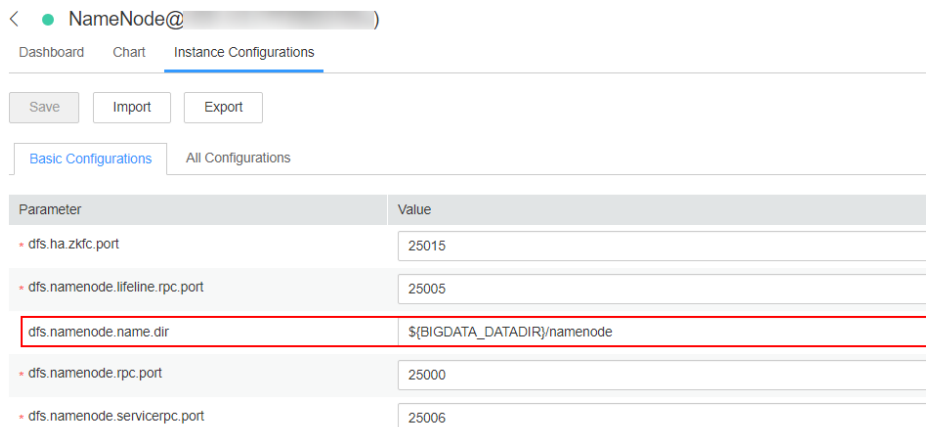
**Step 23** When the standby NameNode fails to push data to the active NameNode as user **omm**, contact the system administrator to handle the fault. Wait for a NameNode metadata combination period and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 24](#).

**Check whether space on the data directory of the active NameNode is insufficient.**

**Step 24** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance**, click the active NameNode of the NameService for which the alarm is generated, and then click **Instance Configurations** to obtain the value of **dfs.namenode.name.dir**. This value is the Fslmage storage directory of the active NameNode.





**Step 25** Log in to the active NameNode as user **root** or **omm**.

**Step 26** Go to the Fslmage storage directory and check the size of the newest Fslmage file (in MB).

**cd** *Storage directory of the active NameNode/current*

**du -m \$(ls -t | grep "fslmage\_[0-9]\*\$" | head -1) | awk '{print \$1}'**

**Step 27** Run the following command to check the available disk space of the active NameNode (in MB).

**df -m ./ | awk 'END{print \$4}'**

**Step 28** Compare the Fslmage file size and the available disk space and determine whether another Fslmage file can be stored on the disk.

- If yes, go to **Step 30**.
- If no, go to **Step 29**.


**Step 29** Clear the redundant files on the disk where the directory resides to reserve sufficient space for metadata. After the clearance, wait for a NameNode metadata combination period and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 30**.

**Collect fault information.**

**Step 30** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 31** Select **NameNode** in the required cluster from the **Service**.

**Step 32** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 33** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.100 ALM-14014 NameNode GC Time Exceeds the Threshold

## Description

The system checks the garbage collection (GC) duration of the NameNode process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14014	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

A long GC duration of the NameNode process may interrupt the services and users cannot read or write files.

## Possible Causes

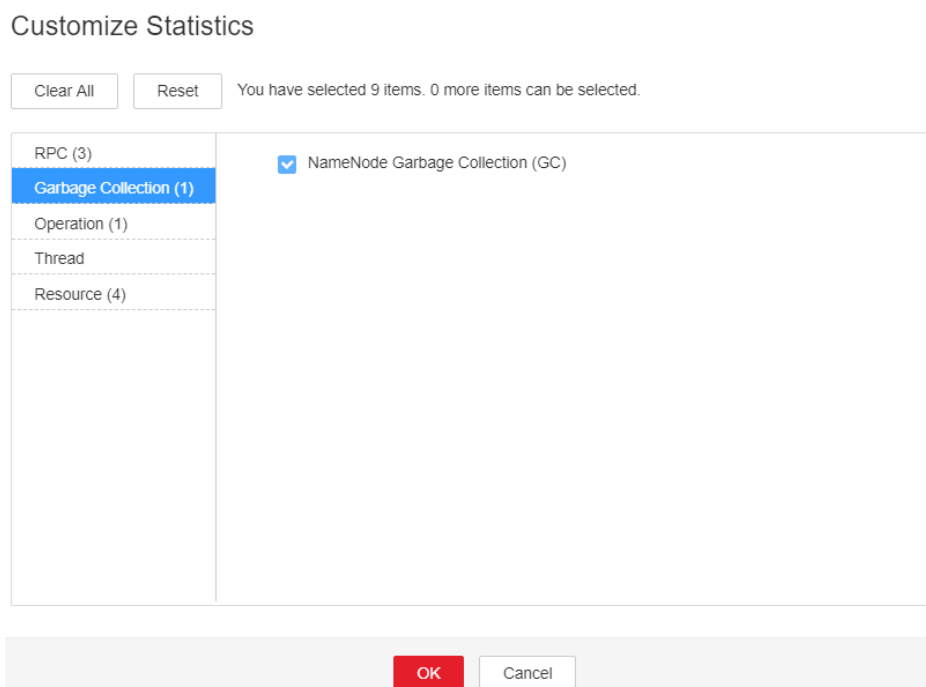
The heap memory of the NameNode instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

## Procedure

### Check the GC duration.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the displayed interface, click the drop-down button of **ALM-14014 NameNode GC Time Exceeds the Threshold**. Then check the role name in **Location** and confirm the IP address of the instance.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Instance > NameNode (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Garbage Collection**, and select **NameNode Garbage Collection (GC)** to check the GC duration statistics of the NameNode process collected every minute.

**Figure 7-71** NameNode Garbage Collection (GC)



- Step 3** Check whether the GC duration of the NameNode process collected every minute exceeds the threshold (12 seconds by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 7**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations > All Configurations > NameNode > System** to increase the value of **GC\_OPTS** parameter as required.

 **NOTE**

The recommended mapping between the number of HDFS file objects (filesystem objects = files + blocks) and the JVM parameters configured for NameNode is as follows:

- If the number of file objects reaches 10,000,000, you are advised to set the JVM parameters as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the number of file objects reaches 20,000,000, you are advised to set the JVM parameters as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G
- If the number of file objects reaches 50,000,000, you are advised to set the JVM parameters as follows: -Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G
- If the number of file objects reaches 100,000,000, you are advised to set the JVM parameters as follows: -Xms64G -Xmx64G -XX:NewSize=6G -XX:MaxNewSize=6G
- If the number of file objects reaches 200,000,000, you are advised to set the JVM parameters as follows: -Xms96G -Xmx96G -XX:NewSize=9G -XX:MaxNewSize=9G
- If the number of file objects reaches 300,000,000, you are advised to set the JVM parameters as follows: -Xms164G -Xmx164G -XX:NewSize=12G -XX:MaxNewSize=12G

**Step 5** Save the configuration and restart the NameNode instance.

 **NOTE**

- During the restart of the active NameNode, a NameNode active/standby switchover occurs. As a result, no active node is available in the system for a short period of time (in the transition period of the active/standby switchover), and an alarm indicating that the HDFS service is unavailable may be generated. In addition, an error is reported during the running read and write tasks, but services are not interrupted.
- Services are not affected after the standby NameNode is restarted.


**Step 6** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **NameNode** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.101 ALM-14015 DataNode GC Time Exceeds the Threshold

### Description

The system checks the garbage collection (GC) duration of the DataNode process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14015	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

### Impact on the System

A long GC duration of the DataNode process may interrupt the services and users cannot read or write files.

### Possible Causes

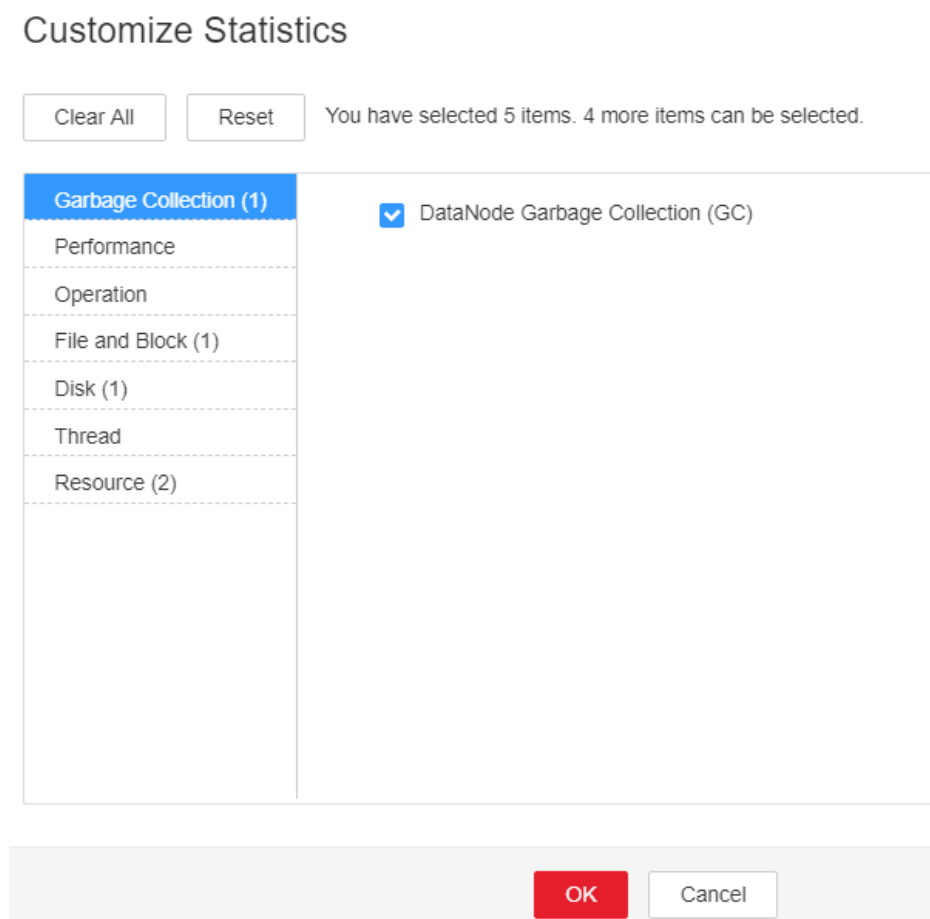
The heap memory of the DataNode instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

## Procedure

### Check the GC duration.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the displayed interface, click the drop-down button of **ALM-14015 DataNode GC Time Exceeds the Threshold**. Then check the role name in **Location** and confirm the IP address of the instance.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Instance > DataNode (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Garbage Collection**, and select **DataNode Garbage Collection (GC)** to check the GC duration statistics of the DataNode process collected every minute.

**Figure 7-72** DataNode Garbage Collection (GC)



- Step 3** Check whether the GC duration of the DataNode process collected every minute exceeds the threshold (12 seconds by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 7**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations > All Configurations > DataNode > System** to increase the value of **GC\_OPTS** parameter as required.

 **NOTE**

The mapping between the average number of blocks of a DataNode instance and the DataNode memory is as follows:

- If the average number of blocks of a DataNode instance reaches 2,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the average number of blocks of a DataNode instance reaches 5,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

**Step 5** Save the configuration and restart the DataNode instance.

 **NOTE**

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.


**Step 6** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **DataNode** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.102 ALM-14016 DataNode Direct Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the direct memory usage of HDFS every 30 seconds. This alarm is generated when the direct memory usage of DataNode instances exceeds the threshold (90% of the maximum memory).

This alarm is automatically cleared when the direct memory usage is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14016	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the available direct memory of DataNode instances is insufficient, a memory overflow may occur and the service breaks down.

## Possible Causes

The direct memory of DataNode instances is overused or the direct memory is inappropriately allocated.

## Handling Procedure

**Check the direct memory usage.**

- Step 1** On the **Home** page of FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the page that is displayed, click the drop-down list in the row containing **ALM-14016 DataNode Direct Memory Usage Exceeds the Threshold**, and view the role name and IP address of the instance for which the alarm is generated in the **Location** area.
- Step 2** On the **Home** page of FusionInsight Manager, choose **Cluster > Services > HDFS**. On the page that is displayed, click the **Instance** tab. In the instance list, select **DataNode** (IP address of the instance for which this alarm is generated). Click the drop-down list in the upper right corner of the chart, choose **Customize > Resource**, and select **DataNode Memory** to check the direct memory usage.



**Step 3** Check whether the used direct memory of a DataNode instance reaches 90% (default threshold) of the maximum direct memory allocated to it.

- If yes, go to [Step 4](#).
- If no, go to [Step 8](#).

**Step 4** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > HDFS > Configurations > All Configurations > DataNode > System**. Check whether **-XX:MaxDirectMemorySize** exists in the **GC\_OPTS** parameter.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** Adjust the value of **-XX:MaxDirectMemorySize**.

1. In **GC\_OPTS**, check the value of **-Xmx** and check whether the node memory is sufficient.

 **NOTE**

You can determine whether the node memory is sufficient based on the actual environment. For example, you can use the following method:

Use the IP address to log in to the instance for which the alarm is generated as user **root** and run the **free -g** command to check the value of **Mem** in the **free** column. The value indicates the available memory of the node. In the following example, the available memory of the node is 4 GB.

```
Mem: total used free shared buff/cache available
..... 112 48 4 10 58 46
```

If the value of **Mem** is at least that of **-Xmx**, the node memory is sufficient. If the value of **Mem** is less than that of **-Xmx**, the node memory is insufficient.

- If yes, change the value of **-XX:MaxDirectMemorySize** to that of **-Xmx**.
- If no, increase **-XX:MaxDirectMemorySize** to a value no larger than that of **Mem**.

2. Save the configuration and restart the DataNode instances.

 **NOTE**

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.

**Step 6** Check whether **ALM-14008 DataNode Heap Memory Usage Exceeds the Threshold** exists.


- If yes, rectify the fault by referring to **ALM-14008 DataNode Heap Memory Usage Exceeds the Threshold**.
- If no, go to [Step 7](#).

**Step 7** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect the fault information.**

**Step 8** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 9** Expand the **Service** drop-down list, and select **DataNode** for the target cluster.
- Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 11** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.103 ALM-14017 NameNode Direct Memory Usage Exceeds the Threshold

### Description

The system checks the direct memory usage of the HDFS service every 30 seconds. This alarm is generated when the direct memory usage of a NameNode instance exceeds the threshold (90% of the maximum memory).

The alarm is cleared when the direct memory usage is less than the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14017	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the available direct memory of the HDFS service is insufficient, a memory overflow occurs and the service breaks down.

## Possible Causes

The direct memory of the NameNode instance is overused or the direct memory is inappropriately allocated.

## Procedure

**Check the direct memory usage.**

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the displayed interface, click the drop-down button of **ALM-14017 NameNode Direct Memory Usage Exceeds the Threshold**. Then check the role name in **Location** and confirm the IP address of the instance.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Instance > NameNode (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Resource**, and select **NameNode Memory** to check the direct memory usage.
- Step 3** Check whether the used direct memory of NameNode reaches 90% of the maximum direct memory specified for NameNode by default.
  - If yes, go to [Step 4](#).
  - If no, go to [Step 8](#).
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations > All Configurations > NameNode > System** to check whether "-XX:MaxDirectMemorySize" exists in the **GC\_OPTS** parameter.
  - If yes, go to [Step 5](#).
  - If no, go to [Step 6](#).
- Step 5** In the **GC\_OPTS** parameter, delete "-XX:MaxDirectMemorySize". Save the configuration and restart the NameNode instance.

 NOTE

- During the restart of the active NameNode, a NameNode active/standby switchover occurs. As a result, no active node is available in the system for a short period of time (in the transition period of the active/standby switchover), and an alarm indicating that the HDFS service is unavailable may be generated. In addition, an error is reported during the running read and write tasks, but services are not interrupted.
- Services are not affected after the standby NameNode is restarted.

**Step 6** Check whether the **ALM-14007 NameNode Heap Memory Usage Exceeds the Threshold** exists.

- If yes, handle the alarm by referring to **ALM-14007 NameNode Heap Memory Usage Exceeds the Threshold**.
- If no, go to [Step 7](#).


**Step 7** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect fault information.**

**Step 8** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 9** Select **NameNode** in the required cluster from the **Service**.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.104 ALM-14018 NameNode Non-heap Memory Usage Exceeds the Threshold

### Description

The system checks the non-heap memory usage of the HDFS NameNode every 30 seconds and compares the actual usage with the threshold. The non-heap memory usage of the HDFS NameNode has a default threshold. This alarm is generated when the non-heap memory usage of the HDFS NameNode exceeds the threshold.

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS** to change the threshold.

This alarm is cleared when the no-heap memory usage of the HDFS NameNode is less than or equal to the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14018	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the memory usage of the HDFS NameNode is too high, data read/write performance of HDFS will be affected.

## Possible Causes

Non-heap memory of the HDFS NameNode is insufficient.

## Procedure

### Delete unnecessary files.

**Step 1** Log in to the HDFS client as user **root**. Run the **cd** command to go to the client installation directory, and run the **source bigdata\_env** command.

If the cluster adopts the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

**Step 2** Run the **hdfs dfs -rm -r file or directory path** command to delete unnecessary files.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check the NameNode JVM non-heap memory usage and configuration.**

**Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**. The HDFS status page is displayed.

**Step 5** In the **Basic Information** area, click **NameNode(Active)**. The HDFS WebUI is displayed.

 **NOTE**

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

**Step 6** On the HDFS WebUI, click the **Overview** tab. In **Summary**, check the numbers of files, directories, and blocks in HDFS.

**Step 7** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** > **All Configurations**. In **Search**, enter **GC\_OPTS** to check the **GC\_OPTS** non-heap memory parameter of **HDFS->NameNode**.

**Adjust system configurations.**

**Step 8** Check whether the non-heap memory is properly configured based on the number of file objects in [Step 6](#) and the non-heap parameters configured for NameNode in [Step 7](#).

- If yes, go to [Step 9](#).
- If no, go to [Step 12](#).

 **NOTE**

The recommended mapping between the number of HDFS file objects (filesystem objects = files + blocks) and the JVM parameters configured for NameNode is as follows:

- If the number of file objects reaches 10,000,000, you are advised to set the JVM parameters as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the number of file objects reaches 20,000,000, you are advised to set the JVM parameters as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G
- If the number of file objects reaches 50,000,000, you are advised to set the JVM parameters as follows: -Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G
- If the number of file objects reaches 100,000,000, you are advised to set the JVM parameters as follows: -Xms64G -Xmx64G -XX:NewSize=6G -XX:MaxNewSize=6G
- If the number of file objects reaches 200,000,000, you are advised to set the JVM parameters as follows: -Xms96G -Xmx96G -XX:NewSize=9G -XX:MaxNewSize=9G
- If the number of file objects reaches 300,000,000, you are advised to set the JVM parameters as follows: -Xms164G -Xmx164G -XX:NewSize=12G -XX:MaxNewSize=12G

**Step 9** Modify the **GC\_OPTS** parameter of the NameNode based on the mapping between the number of file objects and non-heap memory.

**Step 10** Save the configuration and click **Dashboard** > **More** > **Restart Service**.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

**Step 11** Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 12](#).

**Collect fault information.**

**Step 12** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 13** Select the following services in the required cluster from the **Service**.

- ZooKeeper
- HDFS

**Step 14** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 15** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.105 ALM-14019 DataNode Non-heap Memory Usage Exceeds the Threshold

### Description

The system checks the non-heap memory usage of the HDFS DataNode every 30 seconds and compares the actual usage with the threshold. The non-heap memory usage of the HDFS DataNode has a default threshold. This alarm is generated when the non-heap memory usage of the HDFS DataNode exceeds the threshold.

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS** to change the threshold.

This alarm is cleared when the no-heap memory usage of the HDFS DataNode is less than or equal to the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14019	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the memory usage of the HDFS DataNode is too high, data read/write performance of HDFS will be affected.

## Possible Causes

Non-heap memory of the HDFS DataNode is insufficient.

## Procedure

### Delete unnecessary files.

**Step 1** Log in to the HDFS client as user **root**. Run the **cd** command to go to the client installation directory, and run the **source bigdata\_env** command.

If the cluster adopts the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

**Step 2** Run the **hdfs dfs -rm -r file or directory path** command to delete unnecessary files.

**Step 3** Check whether the alarm is cleared.



- If yes, no further action is required.
- If no, go to [Step 4](#).

#### Check the DataNode JVM non-heap memory usage and configuration.

**Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**.

**Step 5** In the **Basic Information** area, click **NameNode(Active)**. The HDFS WebUI is displayed.

#### NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

**Step 6** On the HDFS WebUI, click the **Datanodes** tab to view the number of blocks of all DataNodes that report alarms.

**Step 7** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** > **All Configurations**. In **Search**, enter **GC\_OPTS** to check the **GC\_OPTS** non-heap memory parameter of **HDFS->DataNode**.

#### Adjust system configurations.

**Step 8** Check whether the memory is properly configured based on the number of blocks in [Step 6](#) and the memory parameters configured for DataNode in [Step 7](#).

- If yes, go to [Step 9](#).
- If no, go to [Step 12](#).

#### NOTE

The mapping between the average number of blocks of a DataNode instance and the DataNode memory is as follows:

- If the average number of blocks of a DataNode instance reaches 2,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the average number of blocks of a DataNode instance reaches 5,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

**Step 9** Modify the **GC\_OPTS** parameter of the DataNode based on the mapping between the number of blocks and memory.

**Step 10** Save the configuration and click **Dashboard** > **More** > **Restart Service**.


#### NOTE

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

**Step 11** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

#### Collect fault information.

- Step 12** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 13** Select the following services in the required cluster from the **Service**.
- ZooKeeper
  - HDFS
- Step 14** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 15** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.106 ALM-14020 Number of Entries in the HDFS Directory Exceeds the Threshold

## Description

The system obtains the number of subfiles and subdirectories in a specified directory every hour and checks whether it reaches the percentage of the threshold (the maximum number of subfiles and subdirectories in an HDFS directory, the threshold for triggering an alarm is **90%** by default). If it exceeds the percentage of the threshold, an alarm is triggered.

When the number of subfiles and subdirectories in the directory the alarm is lower than the percentage of the threshold, the alarm is automatically cleared. When the monitoring switch is disabled, alarms corresponding to all directories are cleared. If a directory is removed from the monitoring list, alarms corresponding to the directory are cleared.

### NOTE

- The **dfs.namenode.fs-limits.max-directory-items** parameter specifies the maximum number of subfiles and subdirectories in the HDFS directory. Its default value is **1048576**. If the number of subfiles and subdirectories in a directory exceeds the parameter value, subfiles and subdirectories cannot be created in the directory.
- The **dfs.namenode.directory-items.monitor** parameter specifies the list of directories to be monitored. Its default value is **/tmp,/SparkJobHistory,/mr-history**.
- The **dfs.namenode.directory-items.monitor.enabled** parameter is used to enable or disable the monitoring switch. Its default value is **true**, which means the monitoring switch is enabled by default.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14020	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
NameServiceName	Specifies the NameService service for which the alarm is generated.
Directory	Specifies the directory for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the number of entries in the monitored directory exceeds 90% of the threshold, an alarm is triggered, but entries can be added to the directory. Once the maximum threshold is exceeded, entries will fail to be added to the directory.

## Possible Causes

The number of entries in the monitored directory exceeds 90% of the threshold.

## Procedure

**Check whether unnecessary files exist in the system.**

**Step 1** Log in to the HDFS client as user **root**. Run the **cd** command to go to the client installation directory, and run the **source bigdata\_env** command to set the environment variables.

If the cluster is in security mode, security authentication is required.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

**Step 2** Run the following command to check whether files and directories in the directory with the alarm can be deleted:

```
hdfs dfs -ls Directory with the alarm
```

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

**Step 3** Run the following command to delete unnecessary files.

```
hdfs dfs -rm -r -f File or directory path
```

 **NOTE**

Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.

**Step 4** Wait 1 hour and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Check whether the threshold is correctly configured.**

**Step 5** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** > **All Configurations**. Search for the **dfs.namenode.fs-limits.max-directory-items** parameter and check whether the parameter value is appropriate.

- If yes, go to [Step 9](#).
- If no, go to [Step 6](#).

**Step 6** Increase the parameter value.

**Step 7** Save the configuration and click **Dashboard** > **More** > **Restart Service**.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.


**Step 8** Wait 1 hour and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Collect fault information.**

**Step 9** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

**Step 10** Select **HDFS** in the required cluster from the **Service**.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.107 ALM-14021 NameNode Average RPC Processing Time Exceeds the Threshold

### Description

The system checks the average RPC processing time of NameNode every 30 seconds, and compares the actual average RPC processing time with the threshold (default value: 100 ms). This alarm is generated when the system detects that the average RPC processing time exceeds the threshold for several consecutive times (10 times by default).

You can choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS** to change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the average RPC processing time of NameNode is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the average RPC processing time of NameNode is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14021	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService service for which the alarm is generated.

Name	Meaning
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

NameNode cannot process the RPC requests from HDFS clients, upper-layer services that depend on HDFS, and DataNode in a timely manner. Specifically, the services that access HDFS run slowly or the HDFS service is unavailable.

## Possible Causes

- The CPU performance of NameNode nodes is insufficient and therefore NameNode nodes cannot process messages in a timely manner.
- The configured NameNode memory is too small and frame freezing occurs on the JVM due to frequent full garbage collection.
- NameNode parameters are not configured properly, so NameNode cannot make full use of system performance.

## Procedure

### Obtain alarm information.

**Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.

**Step 2** Check the alarm. Obtain the host name of the NameNode node involved in this alarm from the **HostName** information of **Location**. Then obtain the name of the NameService node involved in this alarm from the **NameServiceName** information of **Location**.

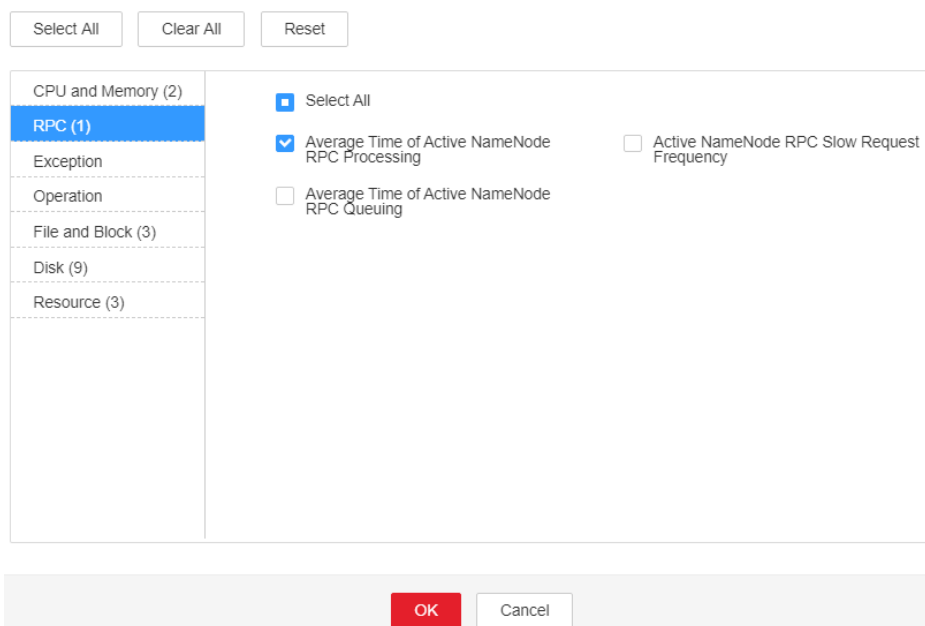
### Check whether the threshold is too small.

**Step 3** Check the status of the services that depend on HDFS. Check whether the services run slowly or task execution times out.

- If yes, go to **Step 8**.
- If no, go to **Step 4**.

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > RPC**, and select **Average Time of Active NameNode RPC Processing** and click **OK**.

**Figure 7-73** Average Time of Active NameNode RPC Processing  
Customize Statistics



**Step 5** On the **Average Time of Active NameNode RPC Processing** monitoring page, obtain the value of the NameService node involved in this alarm.

**Step 6** On the FusionInsight Manager portal, choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS**. Locate **Average Time of Active NameNode RPC Processing** and click the **Modify** in the **Operation** column of the default rule. The **Modify Rule** page is displayed. Change **Threshold** to 150% of the peak value within one day before and after the alarm is generated. Click **OK** to save the new threshold.

**Figure 7-74** Modify Rule

Thresholds > **Modify Rule**

---

\* Rule Name:

\* Alarm Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Others

Thresholds: Start and End Time Threshold

-   ms

**Step 7** Wait for 5 minutes and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Check whether the CPU performance of the NameNode node is sufficient.**

**Step 8** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and check whether **ALM-12016 CPU Usage Exceeds the Threshold** is generated for the NameNode node.

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

**Step 9** Handle **ALM-12016 CPU Usage Exceeds the Threshold** by taking recommended actions.

**Step 10** Wait for 10 minutes and check whether alarm 14021 is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Check whether the memory of the NameNode node is too small.**

**Step 11** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and check whether **ALM-14007 HDFS NameNode Heap Memory Usage Exceeds the Threshold** is generated for the NameNode node.

- If yes, go to [Step 12](#).
- If no, go to [Step 14](#).

**Step 12** Handle **ALM-14007 HDFS NameNode Heap Memory Usage Exceeds the Threshold** by taking recommended actions.



- Step 13** Wait for 10 minutes and check whether alarm 14021 is automatically cleared.
- If yes, no further action is required.
  - If no, go to [Step 14](#).

**Check whether NameNode parameters are configured properly.**


- Step 14** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** > **All Configurations**. Search for parameter **dfs.namenode.handler.count** and view its value. If the value is less than or equal to 128, change it to **128**. If the value is greater than 128 but less than 192, change it to **192**.
- Step 15** Search for parameter **ipc.server.read.threadpool.size** and view its value. If the value is less than 5, change it to **5**.
- Step 16** Click **Save** and click **OK**.
- Step 17** On the **Instance** page of HDFS, select the standby NameNode of NameService involved in this alarm and choose **More** > **Restart Instance**. Enter the password and click **OK**. Wait until the standby NameNode is started up.

 **NOTE**

Services are not affected after the standby NameNode is restarted.

- Step 18** On the **Instance** page of HDFS, select the active NameNode of NameService involved in this alarm and choose **More** > **Restart Instance**. Enter the password and click **OK**. Wait until the active NameNode is started up.
- Step 19** Wait for 1 hour and then check whether the alarm is automatically cleared.
- If yes, no further action is required.
  - If no, go to [Step 20](#).

**Collect fault information.**

- Step 20** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.
- Step 21** Select the following node in the required cluster from the **Service**.
- HDFS
- Step 22** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 23** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.108 ALM-14022 NameNode Average RPC Queuing Time Exceeds the Threshold

### Description

The system checks the average RPC queuing time of NameNode every 30 seconds, and compares the actual average RPC queuing time with the threshold (default value: 200 ms). This alarm is generated when the system detects that the average RPC queuing time exceeds the threshold for several consecutive times (10 times by default).

You can choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS** to change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the average RPC queuing time of NameNode is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the average RPC queuing time of NameNode is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14022	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService service for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

NameNode cannot process the RPC requests from HDFS clients, upper-layer services that depend on HDFS, and DataNode in a timely manner. Specifically, the services that access HDFS run slowly or the HDFS service is unavailable.

## Possible Causes

- The CPU performance of NameNode nodes is insufficient and therefore NameNode nodes cannot process messages in a timely manner.
- The configured NameNode memory is too small and frame freezing occurs on the JVM due to frequent full garbage collection.
- NameNode parameters are not configured properly, so NameNode cannot make full use of system performance.
- The volume of services that access HDFS is too large and therefore NameNode is overloaded.

## Procedure

### Obtain alarm information.

**Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.

**Step 2** Check the alarm. Obtain the alarm generation time from **Generated**. Obtain the host name of the NameNode node involved in this alarm from the **HostName** information of **Location**. Then obtain the name of the NameService node involved in this alarm from the **NameServiceName** information of **Location**.

### Check whether the threshold is too small.

**Step 3** Check the status of the services that depend on HDFS. Check whether the services run slowly or task execution times out.

- If yes, go to [Step 8](#).
- If no, go to [Step 4](#).

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > RPC**, and select **Average Time of Active NameNode RPC Queuing** and click **OK**.

**Figure 7-75** Average Time of Active NameNode RPC Queuing

Customize Statistics

Select All Clear All Reset

CPU and Memory (2)	<input checked="" type="checkbox"/> Select All	<input type="checkbox"/> Average Time of Active NameNode RPC Processing	<input type="checkbox"/> Active NameNode RPC Slow Request Frequency
<b>RPC (1)</b>	<input checked="" type="checkbox"/> Average Time of Active NameNode RPC Queuing		
Exception			
Operation			
File and Block (3)			
Disk (9)			
Resource (3)			

OK Cancel

**Step 5** On the **Average Time of Active NameNode RPC Queuing** monitoring page, obtain the value of the NameService node involved in this alarm.

**Step 6** On the FusionInsight Manager portal, choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS**. Locate **Average Time of Active NameNode RPC Queuing** and click the **Modify** in the **Operation** column of the default rule. The **Modify Rule** page is displayed. Change **Threshold** to 150% of the monitored value. Click **OK** to save the new threshold.

**Step 7** Wait for 1 minute and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Check whether the CPU performance of the NameNode node is sufficient.**

**Step 8** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and check whether **ALM-12016 HDFS NameNode Memory Usage Exceeds the Threshold** is generated.

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

**Step 9** Handle **ALM-12016 CPU Usage Exceeds the Threshold** by taking recommended actions.

**Step 10** Wait for 10 minutes and check whether alarm 14022 is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Check whether the memory of the NameNode node is too small.**

**Step 11** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and check whether **ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold** is generated.

- If yes, go to [Step 12](#).
- If no, go to [Step 14](#).

**Step 12** Handle **ALM-14007 CPU Usage Exceeds the Threshold** by taking recommended actions.

**Step 13** Wait for 10 minutes and check whether alarm 14022 is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 14](#).

**Check whether NameNode parameters are configured properly.**

**Step 14** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations > All Configurations**. Search for parameter **dfs.namenode.handler.count** and view its value. If the value is less than or equal to 128, change it to **128**. If the value is greater than 128 but less than 192, change it to **192**.

**Step 15** Search for parameter **ipc.server.read.threadpool.size** and view its value. If the value is less than 5, change it to 5.

**Step 16** Click **Save**, and click **OK**.

**Step 17** On the **Instance** page of HDFS, select the standby NameNode of NameService involved in this alarm and choose **More > Restart Instance**. Enter the password and click **OK**. Wait until the standby NameNode is started up.

 **NOTE**

Services are not affected after the standby NameNode is restarted.


**Step 18** On the **Instance** page of HDFS, select the active NameNode of NameService involved in this alarm and choose **More > Restart Instance**. Enter the password and click **OK**. Wait until the active NameNode is started up.

**Step 19** Wait for 1 hour and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 20](#).

**Check whether the HDFS workload changes and reduce the workload properly.**

**Step 20** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HDFS**. Click the drop-down menu in the upper right corner of **Chart**, click **Customize**, select **Average Time of Active NameNode RPC Queuing** and click **OK**.

**Figure 7-76** Average Time of Active NameNode RPC Queuing

**Step 21** Click . The **Details** page is displayed.

**Step 22** Set the monitoring data display period, from 5 days before the alarm generation time to the alarm generation time. Click **OK**.

**Step 23** On the **Average RPC Queuing Time** monitoring page, check whether the point in time when the queuing time increases abruptly exists.

- If yes, go to **Step 24**.
- If no, go to **Step 27**.

**Step 24** Confirm and check the point in time. Check whether a new task frequently accesses HDFS and whether the access frequency can be reduced.

**Step 25** If a Balancer task starts at the point in time, stop the task or specify a node for the task to reduce the HDFS workload.

**Step 26** Wait for 1 hour and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 27**.

#### Collect fault information.

**Step 27** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 28** Select **HDFS** in the required cluster from the **Service**.

**Step 29** Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 30** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.109 ALM-14023 Percentage of Total Reserved Disk Space for Replicas Exceeds the Threshold

### Description

The system checks the percentage of total reserved disk space for replicas (Total reserved disk space for replicas/(Total reserved disk space for replicas + Total remaining disk space)) every 30 seconds and compares the actual percentage with the threshold (**90%** by default). This alarm is generated when the percentage of total reserved disk space for replicas exceeds the threshold for multiple consecutive times (**Trigger Count**).

The alarm is cleared in the following two scenarios: The value of **Trigger Count** is **1** and the percentage of total reserved disk space for replicas is less than or equal to the threshold; the value of **Trigger Count** is greater than **1** and the percentage of total reserved disk space for replicas is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14023	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
NameServiceName	Specifies the NameService service for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

The performance of writing data to HDFS is affected. If all remaining DataNode space is reserved for replicas, writing HDFS data fails.

## Possible Causes

- The alarm threshold is improperly configured.
- The disk space configured for the HDFS cluster is insufficient.
- The volume of services that access HDFS is too large and therefore DataNode is overloaded.

## Procedure


Check whether the alarm threshold is appropriate.

**Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS > Disk > Percentage of Reserved Space for Replicas of Unused Space** to check whether the alarm threshold is appropriate. (The default threshold is **90%**. Users can change it as required.)

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

**Step 2** Choose **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS > Disk > Percentage of Reserved Space for Replicas of Unused Space** and Click **Modify**, change the threshold based on the actual usage.

Figure 7-77 Modify Thresholds



**Step 3** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Check whether an alarm indicating insufficient disk space is generated.

**Step 4** On the FusionInsight Manager portal, check whether **ALM-14001 HDFS Disk Usage Exceeds the Threshold** or **ALM-14002 DataNode Disk Usage Exceeds the Threshold** exists on the **O&M > Alarm > Alarms** page.

- If yes, go to [Step 5](#).
- If no, go to [Step 7](#).

**Step 5** Handle the alarm by referring to instructions in **ALM-14001 HDFS Disk Usage Exceeds the Threshold** or **ALM-14002 DataNode Disk Usage Exceeds the Threshold** and check whether the alarm is cleared.



- If yes, go to [Step 6](#).
- If no, go to [Step 7](#).

**Step 6** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Expand the DataNode capacity.**

**Step 7** Expand the DataNode capacity.


**Step 8** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Collect fault information.**

**Step 9** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 10** Select **HDFS** in the required cluster from the **Service**.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.110 ALM-14024 Tenant Space Usage Exceeds the Threshold

### Description

The system checks the space usage (used space of each directory/space allocated to each directory) of each directory associated with a tenant every hour and compares the space usage of each directory with the threshold set for the directory. This alarm is generated when the space usage exceeds the threshold.

This alarm is cleared when the space usage is less than or equal to the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14024	Minor	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TenantName	Specifies the tenant for which the alarm is generated.
DirectoryName	Specifies the directory for which the alarm is generated.
Trigger condition	Specifies the threshold for triggering the alarm.

## Impact on the System

This alarm is generated if the space usage of the tenant directory exceeds the custom threshold. File writing to the directory is not affected. If the used space exceeds the maximum storage space allocated to the directory, the HDFS fails to write data to the directory.

## Possible Causes

- The alarm threshold is improperly configured.
- The space allocated to the tenant is improper.

## Procedure

**Check whether the alarm threshold is appropriate.**

- Step 1** View the alarm location information to obtain the tenant name and tenant directory for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose the **Tenant Resources** page, select the tenant for which the alarm is generated, and click **Resources**. Check whether

the storage space threshold configured for the tenant directory for which the alarm is generated is proper. (The default value 90% is a proper value. You can set it based on the site requirements.)

- If yes, go to [Step 5](#).
- If no, go to [Step 3](#).

**Step 3** On the **Resources** page, click **Modify** to modify or delete the storage space threshold.

**Step 4** About one minute later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Check whether the space allocated to the tenant is appropriate.**

**Step 5** On the FusionInsight Manager portal, choose the **Tenant Resources** page, select the tenant for which the alarm is generated, and click **Resources**. Check whether the storage space quota of the tenant directory for which the alarm is generated is proper based on the actual service status of the tenant directory.

- If yes, go to [Step 8](#).
- If no, go to [Step 6](#).

**Step 6** On the **Resources** page, click **Modify** to modify the storage space quota.


**Step 7** About one minute later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect fault information.**

**Step 8** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 9** Select **HDFS** in the required cluster and **NodeAgent** under **Manager** from the **Service**.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.111 ALM-14025 Tenant File Object Usage Exceeds the Threshold

### Description

The system checks the file object usage (used file objects of each directory/ number of file objects allocated to each directory) of each directory associated with a tenant every hour and compares the file object usage of each directory with the threshold set for the directory. This alarm is generated when the file object usage exceeds the threshold.

This alarm is cleared when the file object usage is less than or equal to the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14025	Minor	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TenantName	Specifies the tenant for which the alarm is generated.
DirectoryName	Specifies the directory for which the alarm is generated.
Trigger condition	Specifies the threshold for triggering the alarm.

### Impact on the System

This alarm is generated if the usage of file objects in a tenant directory exceeds the custom threshold. File writing to the directory is not affected. If the number of used file objects exceeds the maximum number of file objects allocated to the directory, the HDFS fails to write data to the directory.

## Possible Causes

- The alarm threshold is improperly configured.
- The maximum number of file objects allocated to the tenant directory is inappropriate.

## Procedure

### Check whether the alarm threshold is appropriate.


- Step 1** View the alarm location information to obtain the tenant name and tenant directory for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose the **Tenant Resources** page, select the tenant for which the alarm is generated, and click **Resources**. Check whether the file object threshold configured for the tenant directory for which the alarm is generated is proper. (The default value 90% is a proper value. You can set it based on the site requirements.)
- If yes, go to [Step 5](#).
  - If no, go to [Step 3](#).
- Step 3** On the **Resources** page, click **Modify** to modify or delete the file object threshold of the tenant directory for which the alarm is generated.
- Step 4** About one minute later, check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 5](#).

### Check whether the maximum number of file objects allocated to the tenant is appropriate.

- Step 5** On the FusionInsight Manager portal, choose the **Tenant Resources** page, select the tenant for which the alarm is generated, and click **Resources**. Check whether the maximum number of file objects configured for the tenant directory for which the alarm is generated is proper based on the actual service status of the tenant directory.
- If yes, go to [Step 8](#).
  - If no, go to [Step 6](#).
- Step 6** On the **Resources** page, click **Modify** to modify or delete the maximum number of file objects configured for the tenant directory.
- Step 7** About one minute later, check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 8](#).

### Collect fault information.

- Step 8** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 9** Select **HDFS** in the required cluster and **NodeAgent** under **Manager** from the **Service**.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.112 ALM-14026 Blocks on DataNode Exceed the Threshold

## Alarm Description

The system checks the number of blocks on each DataNode every 30 seconds. This alarm is generated when the number of blocks on the DataNode exceeds the threshold.

If **Trigger Count** is **1** and the number of blocks on the DataNode is less than or equal to the threshold, this alarm is cleared. If **Trigger Count** is greater than **1** and the number of blocks on the DataNode is less than or equal to 90% of the threshold, this alarm is cleared.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14026	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Parameter	Description
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If this alarm is reported, there are too many blocks on the DataNode. In this case, data writing into the HDFS may fail due to insufficient disk space.

## Possible Causes

- The alarm threshold is improperly configured.
- Data skew occurs among DataNodes.
- The disk space configured for the HDFS cluster is insufficient.

## Handling Procedure

### Change the threshold.

- Step 1** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **HDFS**. Then choose **Configurations > All Configurations**. On the displayed page, find the **GC\_OPTS** parameter under **HDFS->DataNode**.
- Step 2** Set the threshold of the DataNode blocks. Specifically, change the value of **Xmx** of the **GC\_OPTS** parameter. **Xmx** specifies the memory, and each GB memory supports a maximum of 500,000 DataNode blocks. Set the memory as required. Confirm that **GC\_PROFILE** is set to **custom** and save the configuration.
- Step 3** Choose **Cluster**, click the name of the desired cluster, and choose **HDFS > Instance**. Select the DataNode instance whose status is **Expired**, click **More**, and select **Restart Instance** to make the **GC\_OPTS** configuration take effect.

### NOTE

Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.

- Step 4** Check whether the alarm is cleared 5 minutes later.
- If yes, no further action is required.
  - If no, go to [Step 5](#).

### Check whether associated alarms are reported.

- Step 5** On FusionInsight Manager, choose **O&M > Alarm > Alarms**, and check whether the **ALM-14002 DataNode Disk Usage Exceeds the Threshold** alarm exists.
- If yes, go to [Step 6](#).
  - If no, go to [Step 8](#).
- Step 6** Handle the alarm by following the instructions in **ALM-14002 DataNode Disk Usage Exceeds the Threshold** and check whether the alarm is cleared.
- If yes, go to [Step 7](#).

- If no, go to [Step 8](#).

**Step 7** Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Expand the DataNode capacity.**

**Step 8** Expand the DataNode capacity.


**Step 9** On FusionInsight Manager, wait for 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

**Collect fault information.**

**Step 10** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 11** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

**Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

### Configuration rules of the DataNode JVM parameter.

Default value of the DataNode JVM parameter **GC\_OPTS**:

```
-Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M -
XX:MetaspaceSize=128M -XX:MaxMetaspaceSize=128M -
XX:+UseConcMarkSweepGC -XX:+CMSParallelRemarkEnabled -
XX:CMSInitiatingOccupancyFraction=65 -XX:+PrintGCDetails -
Dsun.rmi.dgc.client.gcInterval=0x7FFFFFFFFFFFFFFFE -
Dsun.rmi.dgc.server.gcInterval=0x7FFFFFFFFFFFFFFFE -XX:-
OmitStackTraceInFastThrow -XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation
-XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=1M -
Djdk.tls.ephemeralDHKeySize=2048
```

Average number of saved blocks = Number of HDFS blocks x 3/Number of DataNodes



 NOTE

To obtain the number of HDFS blocks, log in to FusionInsight Manager, choose **Cluster > Services > HDFS**, click **NameNode(xxx,Active)** to the right of **NameNode Web UI** to go to the native HDFS web UI, and then view the information in the **Summary** area.

## Summary

Security is on.  
 Safemode is off.  
 1,580 files and directories, 1,183 blocks (1,183 replicated blocks, 0 erasure coded block groups) = 2,763 total filesystem object(s).  
 Heap Memory used 179.06 MB of 1.99 GB Heap Memory. Max Heap Memory is 3.98 GB.  
 Non Heap Memory used 134 MB of 137.06 MB Committed Non Heap Memory. Max Non Heap Memory is 488 MB.

If the average number of blocks on a single DataNode has changed, modify -**Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M** in the default value. The following table lists the reference values.

**Table 7-102** DataNode JVM configuration

Average Number of Blocks in a DataNode Instance	Reference Value
2,000,000	-Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
5,000,000	-Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

**Xmx** specifies memory which corresponds to the threshold of the number of DataNode blocks, and each GB memory supports a maximum of 500,000 DataNode blocks. Set the memory as required.

## 7.12.113 ALM-14027 DataNode Disk Fault

### Alarm Description

The system checks the disk status on DataNodes every 60 seconds. This alarm is generated when a disk is faulty.

After all faulty disks on the DataNode are recovered, you need to manually clear the alarm and restart the DataNode.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14027	Major	No

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Failed Volumes	Specifies the list of faulty disks.

## Impact on the System

If this alarm is reported, there are abnormal disk partitions on the DataNode. This may cause the loss of written files.

## Possible Causes

- The hard disk is faulty.
- The disk permissions are configured improperly.

## Handling Procedure

**Check whether a disk alarm is generated.**

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms** and check whether **ALM-12014 Partition Lost** or **ALM-12033 Slow Disk Fault** exists.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

**Step 2** Rectify the fault by referring to the handling procedure of **ALM-12014 Partition Lost** or **ALM-12033 Slow Disk Fault**. Then, check whether the alarm is cleared.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

**Step 3** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

**Modify disk permissions.**

**Step 4** Choose **O&M > Alarm > Alarms** and view **Location** and **Additional Information** of the alarm to obtain the location of the faulty disk.

**Step 5** Log in to the node for which the alarm is generated as user **root**. Go to the directory where the faulty disk is located, and run the **ll** command to check whether the permission of the faulty disk is **711** and whether the user is **omm**.

- If yes, go to [Step 8](#).
- If no, go to [Step 6](#).

**Step 6** Modify the permission of the faulty disk. For example, if the faulty disk is **data1**, run the following commands:

```
chown omm:wheel data1
```

```
chmod 711 data1
```

**Step 7** In the alarm list on Manager, click **Clear** in the **Operation** column of the alarm to manually clear the alarm. Choose **Cluster > Services > HDFS > Instance**, select the DataNode, choose **More > Restart Instance**, wait for 5 minutes, and check whether a new alarm is reported.

 **NOTE**


Services may be affected or interrupted during the restart. You are advised to perform this operation during off-peak hours.

- If no, no further action is required.
- If yes, go to [Step 8](#).

**Collect the fault information.**

**Step 8** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 9** Expand the **Service** drop-down list, and select **HDFS** and **OMS** for the target cluster.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

After the fault is rectified, the system does not automatically clear this alarm and you need to manually clear the alarm.

## Related Information

None

## 7.12.114 ALM-14028 Number of Blocks to Be Supplemented Exceeds the Threshold

### Alarm Description

The system checks the number of blocks to be supplemented every 30 seconds and compares the number with the threshold. The number of blocks to be supplemented has a default threshold. This alarm is generated when the number of blocks to be supplemented exceeds the threshold.

You can change the threshold specified by **Blocks Under Replicated (NameNode)** by choosing **O&M > Alarm > Thresholds > Name of the desired cluster > HDFS > File and Block**.

If **Trigger Count** is set to **1** and the number of blocks to be supplemented is less than or equal to the threshold, this alarm is cleared. If **Trigger Count** is greater than **1** and the number of blocks to be supplemented is less than or equal to 90% of the threshold, this alarm is cleared.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14028	Minor	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

Data stored in HDFS is lost. HDFS may enter the security mode and cannot provide write services. Lost block data cannot be restored.

## Possible Causes

- The DataNode instance is abnormal.
- Data is deleted.
- The number of replicas written into the file is greater than the number of DataNodes.

## Handling Procedure

- Step 1** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, check whether alarm **ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold** is generated.
- If yes, go to **Step 2**.
  - If no, go to **Step 3**.
- Step 2** Rectify the fault according to the handling procedure of **ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold**. Five minutes later, check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 3**.
- Step 3** Log in to the HDFS client as user **root**. The user password is defined by the user before the installation. Contact the MRS cluster administrator to obtain the password. Run the following commands:
- Security mode:  
`cd Client installation directory`  
`source bigdata_env`  
`kinit hdfs`
  - Normal mode:  
`su - omm`  
`cd Client installation directory`  
`source bigdata_env`
- Step 4** Run the `hdfs fsck / >> fsck.log` command to obtain the status of the current cluster.
- Step 5** Run the following command to count the number ( $M$ ) of blocks to be replicated:
- ```
cat fsck.log | grep "Under-replicated"
```
- Step 6** Run the following command to count the number (N) of blocks to be replicated in the `/tmp/hadoop-yarn/staging/` directory:
- ```
cat fsck.log | grep "Under replicated" | grep "/tmp/hadoop-yarn/staging/" | wc -l
```

### NOTE

`/tmp/hadoop-yarn/staging/` is the default directory. If the directory is modified, obtain it from the configuration item `yarn.app.mapreduce.am.staging-dir` in the `mapred-site.xml` file.

**Step 7** Check whether the percentage of  $N$  is greater than 50% ( $N/M > 50\%$ ).

- If yes, go to [Step 8](#).
- If no, go to [Step 9](#).

**Step 8** Run the following command to reconfigure the number of file replicas in the directory (set the number of file replicas to the number of DataNodes or the default number of file replicas):

```
hdfs dfs -setrep -w Number of file replicas/tmp/hadoop-yarn/staging/
```

 **NOTE**

To obtain the default number of file replicas:

Log in to FusionInsight Manager, choose **Cluster > Services > HDFS > Configurations > All Configurations**, and search for the **dfs.replication** parameter. The value of this parameter is the default number of file replicas.


Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Collect the fault information.**

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.115 ALM-14029 Number of Blocks in a Replica Exceeds the Threshold

## Alarm Description

The system checks the number of blocks in a single replica every four hours and compares the number with the threshold. There is a threshold for the number of blocks in a single replica. This alarm is generated when the actual number of blocks in a single replica exceeds the threshold.

This alarm is cleared when the number of blocks to be supplemented is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14029	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Replica data is prone to be lost when a node is faulty. Too many files of a single replica affect the security of the HDFS file system.

## Possible Causes

- The DataNode is faulty.
- The disk is faulty.
- Files are written to a single replica.

## Handling Procedure

- Step 1** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, check whether alarm **ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold** is generated.
  - If yes, go to **Step 2**.
  - If no, go to **Step 3**.
- Step 2** Rectify the fault according to the handling procedure of **ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold**. In the next detection period, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

**Step 3** Check whether files of a single replica have been written into the service.

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

**Step 4** Log in to the HDFS client as user **root**. The user password is defined by the user before the installation. Contact the MRS cluster administrator to obtain the password. Run the following commands:

- Security mode:  
`cd Client installation directory`  
`source bigdata_env`  
`kinit hdfs`
- Normal mode:  
`su - omm`  
`cd Client installation directory`  
`source bigdata_env`

**Step 5** Run the following command on the client node to increase the number of replicas for a single replica file:

```
hdfs dfs -setrep -w file replica number file name or file path
```


**Step 6** In the next detection period, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect the fault information.**

**Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 8** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None



## 7.12.116 ALM-14030 HDFS Allows Write of Single-Replica Data

### Alarm Description

This alarm is generated when **dfs.single.replication.enable** is set to **true**, indicating that HDFS is configured to allow write of single-replica data.

This alarm is cleared when this function is disabled on HDFS.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14030	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

### Impact on the System

If this configuration is enabled on the server and the number of HDFS replicas configured on the client is 1, single-replica data can be written to HDFS. Data of a single replica may be lost. Therefore, the system does not allow write of single-replica data by default. If a service requires single-replica data write to a directory, modify the HDFS configuration item **dfs.single.replication.exclude.pattern**.

### Possible Causes

The HDFS configuration item **dfs.single.replication.enable** is set to **true**.

### Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **Cluster > Services > HDFS**. On the page that is displayed, click the **Configurations** tab then the **All Configurations** sub-tab.
- Step 2** Search for **dfs.single.replication.enable** in the search box, change the value of the configuration item to **false**, and click **Save**.


**Step 3** Wait for about 10 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 5** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

**Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 7** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.117 ALM-14031 DataNode Process Is Abnormal

### Alarm Description

The DataNode process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14031	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

## Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

## Handling Procedure

**Check whether the process is in the D, Z, or T state.**

**Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.

- If the alarm is not in the list, no further action is required.
- If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in [Step 2](#).

**Step 2** Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.

**Step 3** Run the following command to check the process state:

```
ps ww -eo stat,cmd| grep -w
org.apache.hadoop.hdfs.server.datanode.DataNode | grep -v grep | awk '{print
$1}'
```

**Step 4** Check whether the command output contains any abnormal state (D, Z, or T).

- If the output contains any abnormal state, go to [Step 5](#).
- If the output does not contain abnormal states, go to [Step 7](#).

**Step 5** Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)

**Step 6** Wait 5 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to [Step 7](#).

**Collect fault information.**

- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 8** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.
- Step 9** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.

----End

**Alarm Clearance**

This alarm is automatically cleared after the fault is rectified.

**Related Information**

None.

**7.12.118 ALM-14032 JournalNode Process Is Abnormal**

**Alarm Description**

The JournalNode process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

**Alarm Attributes**

Alarm ID	Alarm Severity	Auto Cleared
14032	Major	Yes

**Alarm Parameters**

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

## Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

## Handling Procedure

**Check whether the process is in the D, Z, or T state.**

**Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.

- If the alarm is not in the list, no further action is required.
- If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in [Step 2](#).

**Step 2** Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.

**Step 3** Run the following command to check the process state:

```
ps ww -eo stat,cmd| grep -w
org.apache.hadoop.hdfs.qjournal.server.JournalNode | grep -v grep | awk
'{print$1}'
```

**Step 4** Check whether the command output contains any abnormal state (D, Z, or T).

- If the output contains any abnormal state, go to [Step 5](#).
- If the output does not contain abnormal states, go to [Step 7](#).

**Step 5** Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)

**Step 6** Wait 5 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to [Step 7](#).

**Collect fault information.**

**Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 8** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.
- Step 9** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.119 ALM-14033 ZKFC Process Is Abnormal

### Alarm Description

The ZKFC process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14033	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

## Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

## Handling Procedure

**Check whether the process is in the D, Z, or T state.**

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.
- If the alarm is not in the list, no further action is required.
  - If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in [Step 2](#).
- Step 2** Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.
- Step 3** Run the following command to check whether the process state is abnormal:
- ```
ps ww -eo stat,cmd| grep -w  
org.apache.hadoop.hdfs.tools.DFSZKFailoverController | grep -v grep | awk  
'{print$1}'
```
- Step 4** Check whether the command output contains any abnormal state (D, Z, or T).
- If the output contains any abnormal state, go to [Step 5](#).
 - If the output does not contain abnormal states, go to [Step 7](#).
- Step 5** Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)
- Step 6** Wait 5 minutes and check whether the alarm is cleared.
- If the alarm is cleared, no further action is required.
 - If the alarm fails to be cleared, go to [Step 7](#).
- Collect fault information.**
- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 8** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.
- Step 9** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.120 ALM-14034 Router Process Is Abnormal

Alarm Description

The Router process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 14034 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

Handling Procedure

Check whether the process is in the D, Z, or T state.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.
- If the alarm is not in the list, no further action is required.
 - If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in [Step 2](#).
- Step 2** Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.
- Step 3** Run the following command to check whether the process state is abnormal:
- ```
ps ww -eo stat,cmd| grep -w org.apache.hadoop.hdfs.server.federation.router.DFSRouter | grep -v grep | awk '{print$1}'
```
- Step 4** Check whether the command output contains any abnormal state (D, Z, or T).
- If the output contains any abnormal state, go to [Step 5](#).
  - If the output does not contain abnormal states, go to [Step 7](#).
- Step 5** Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)
- Step 6** Wait 5 minutes and check whether the alarm is cleared.
- If the alarm is cleared, no further action is required.
  - If the alarm fails to be cleared, go to [Step 7](#).
- Collect fault information.**
- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 8** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.
- Step 9** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.121 ALM-14035 HttpFS Process Is Abnormal

### Alarm Description

The HttpFS process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14035	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

### Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

### Handling Procedure

**Check whether the process is in the D, Z, or T state.**

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.
- If the alarm is not in the list, no further action is required.
  - If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in [Step 2](#).
- Step 2** Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.
- Step 3** Run the following command to check whether the process state is abnormal:
- ```
ps ww -eo stat,cmd| grep -w  
org.apache.hadoop.fs.http.server.HttpFSServerWebServer | grep -v grep | awk  
'{print$1}'
```
- Step 4** Check whether the command output contains any abnormal state (D, Z, or T).
- If the output contains any abnormal state, go to [Step 5](#).
 - If the output does not contain abnormal states, go to [Step 7](#).
- Step 5** Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)
- Step 6** Wait 5 minutes and check whether the alarm is cleared.
- If the alarm is cleared, no further action is required.
 - If the alarm fails to be cleared, go to [Step 7](#).
- Collect fault information.**
- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 8** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.
- Step 9** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.122 ALM-16000 Percentage of Sessions Connected to the HiveServer to Maximum Number Allowed Exceeds the Threshold

Description

The system detects the percentage of sessions connected to the HiveServer to the maximum number of allowed sessions every 30 seconds. This indicator can be viewed on the **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** > *HiveServer instance*. This alarm is generated when the percentage exceeds the default value **90%**.

To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Hive** > **Percentage of Sessions Connected to the HiveServer to Maximum Number of Sessions Allowed by the HiveServer**.

When the **Trigger Count** is 1, this alarm is cleared when the percentage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the percentage is less than or equal to 90% of the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 16000 | Minor | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

If a connection number alarm is generated, the number of sessions connected to HiveServer is too large. As a result, new connections cannot be established, new tasks fail, or even services restart unexpectedly.

Possible Causes

Too many clients are connected to HiveServer.

Procedure

Increase the maximum number of connections to Hive.

Step 1 On the FusionInsight Manager portal, Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations**.

Step 2 Search for **hive.server.session.control.maxconnections** and increase the value of this parameter. If the value of this parameter is **A**, the threshold is **B**, and the number of sessions connected to the HiveServer is **C**, adjust the value of this parameter according to $A \times B > C$. To view the number of sessions connected to the HiveServer, check the value of **Statistics for Sessions of the HiveServer** on the Hive monitoring page.


Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Collect fault information.

Step 4 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 5 Select **Hive** in the required cluster from the **Service**.

Step 6 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.123 ALM-16001 Hive Warehouse Space Usage Exceeds the Threshold

Description

This alarm is generated when the Hive warehouse space usage exceeds the specified threshold (85% by default). The system checks the Hive data warehouse space usage every 30s. The indicator **Percentage of HDFS Space Used by Hive to the Available Space** can be viewed on the Hive service monitoring page.

To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Hive > Percentage of HDFS Space Used by Hive to the Available Space**.

When the **Trigger Count** is 1, this alarm is cleared when the Hive warehouse space usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the Hive warehouse space usage is less than or equal to 90% of the threshold.

NOTE

The administrator can reduce the warehouse space usage by expanding the warehouse capacity or releasing the used space.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 16001 | Minor | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

The system cannot write data properly. Some data may be lost.

Possible Causes

- The upper limit of the HDFS capacity available for Hive is too small.
- The HDFS space is insufficient.
- Some data nodes break down.

Procedure

Expand the system configuration.

- Step 1** Analyze the cluster HDFS capacity usage and increase the upper limit of the HDFS capacity available for Hive.

Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations**, find **hive.metastore.warehouse.size.percent**, and increase its value so that larger HDFS capacity will be available for Hive. Assume that the value of the configuration item is A, the total HDFS storage space is B, the threshold is C, and the HDFS space used by Hive is D. The adjustment policy is $A \times B \times C > D$. The total HDFS storage space can be viewed on the HDFS NameNode page. The HDFS space used by Hive can be viewed on the Hive monitoring page.

- Step 2** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 3](#).

Expand the system.

- Step 3** Expand the system.

- Step 4** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 5](#).

Check whether the data node is normal.

- Step 5** On the FusionInsight Manager portal, click **O&M** > **Alarm** > **Alarms**.

- Step 6** Check whether "ALM-12006 Node Fault", "ALM-12007 Process Fault", or "ALM-14002 DataNode Disk Usage Exceeds the Threshold" exist.
- If yes, go to [Step 7](#).
 - If no, go to [Step 9](#).

- Step 7** Clear the alarm by following the steps provided in "ALM-12006 Node Fault", "ALM-12007 Process Fault", and "ALM-14002 DataNode Disk Usage Exceeds the Threshold".


- Step 8** Check whether the alarm is cleared.
- If yes, no further action is required.

- If no, go to [Step 9](#).

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 10 Select **Hive** in the required cluster from the **Service**.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.124 ALM-16002 Hive SQL Execution Success Rate Is Lower Than the Threshold

Description

The system checks the percentage of the HQL statements that are executed successfully in every 30 seconds. The formula is: Percentage of HQL statements that are executed successfully = Number of HQL statements that are executed successfully by Hive in a specified period/Total number of HQL statements that are executed by Hive. This indicator can be viewed on the **Cluster > Name of the desired cluster > Services > Hive > Instance > HiveServer instance**. The default threshold of the percentage of HQL statements that are executed successfully is **90%**. An alarm is reported when the percentage is lower than the **90%**. Users can view the name of the host where an alarm is generated in the location information about the alarm. The IP address of the host is the IP address of the HiveServer node.

Users can modify the threshold by choosing **O&M > Alarm > Thresholds > Name of the desired cluster > Hive > Percentage of HQL Statements That Are Executed Successfully by Hive**.

This alarm is cleared when the execution success rate is higher than 110% of the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 16002 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

The service execution capability of the system is too low and cannot properly respond to customer requests. The Hive service is not affected. You need to check HiveServer logs to locate the SQL failure cause.

Possible Causes

- A syntax error occurs in HQL statements.
- The HBase service is abnormal when a Hive on HBase task is performed.
- The Spark service is abnormal when a Hive on Spark task is performed.
- The dependent basic services, such as HDFS, Yarn, and ZooKeeper, are abnormal.

Procedure

Check whether the HQL statements comply with syntax.

- Step 1** On the FusionInsight Manager page, choose **O&M > Alarm** to view the alarm details and obtain the node where the alarm is generated.
- Step 2** Use the Hive client to log in to the HiveServer node where an alarm is reported. Query the HQL syntax provided by Apache, and check whether the HQL commands are correct. For details, see <https://cwiki.apache.org/confluence/display/hive/languagemanual>.
- If yes, go to **Step 4**.
 - If no, go to **Step 3**.

 NOTE

To view the user who runs an incorrect statement, you can download the hiveserver audit log file of the HiveServer node where this alarm is generated. **Start Data** and **End Data** are 10 minutes before and after the alarm generation time respectively. Open the log file and search for the **Result=FAIL** keyword to filter the log information about the incorrect statement, and then view the user who runs the incorrect statement according to **UserName** in the log information.

Step 3 Enter the correct HQL statements, and check whether the command can be properly executed.

- If yes, go to [Step 12](#).
- If no, go to [Step 4](#).

Check whether the HBase service is abnormal.

Step 4 Check whether an Hive on HBase task is performed with the user who runs the HQL command.

- If yes, go to [Step 5](#).
- If no, go to [Step 8](#).

Step 5 On the FusionInsight Manager page, click **Cluster** > *Name of the desired cluster* > **Services**, check whether the HBase service is normal in the service list.

- If yes, go to [Step 8](#).
- If no, go to [Step 6](#).

Step 6 Choose **O&M** > **Alarm**, check the related alarms displayed on the alarm page and clear them according to related alarm help.

Step 7 Enter the correct HQL statements, and check whether the command can be properly executed.

- If yes, go to [Step 12](#).
- If no, go to [Step 8](#).

Check whether the HDFS, Yarn, and ZooKeeper are normal.

Step 8 On the FusionInsight Manager portal, click **Cluster** > *Name of the desired cluster* > **Services**.

Step 9 In the service list, check whether the services, such as HDFS, Yarn, and ZooKeeper are normal.

- If yes, go to [Step 12](#).
- If no, go to [Step 10](#).

Step 10 Check the related alarms displayed on the alarm page and clear them according to related alarm help.

Step 11 Enter the correct HQL statements, and check whether the command can be properly executed.

- If yes, go to [Step 12](#).
- If no, go to [Step 13](#).

Step 12 After 1 minute, check whether the alarm is cleared.

- If yes, no further action is required.


- If no, go to [Step 13](#).

Collect fault information.

Step 13 On the FusionInsight Manager home page, choose **O&M > Log > Download**.

Step 14 Select the following nodes in the required cluster from the **Service**:

- MapReduce
- Hive

Step 15 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.125 ALM-16003 Background Thread Usage Exceeds the Threshold

Description

The system checks the background thread usage in every 30 seconds. This alarm is generated when the usage of the background thread pool of Hive exceeds the threshold, 90% by default.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16003 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |

| Name | Meaning |
|-------------------|--|
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

There are too many background threads, so the newly submitted task cannot run in time.

Possible Causes

The usage of the background thread pool of Hive is excessively high when:

- There are many tasks executed in the background thread pool of HiveServer.
- The capacity of the background thread pool of HiveServer is too small.

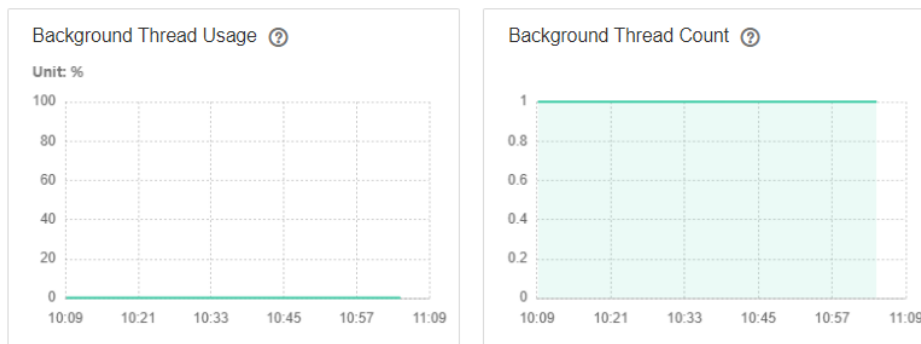
Procedure

Check the number of tasks executed in the background thread pool of HiveServer.

- Step 1** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive**. On the displayed page, click **HiveServer Instance** and check values of **Background Thread Count** and **Background Thread Usage**.

Figure 7-78 Background

Chart



- Step 2** Check whether the number of background threads in the latest half an hour is excessively high. (By default, the queue number is 100, and the thread number is considered as high if it is 90 or larger.)

- If it is, go to [Step 3](#).
- If it is not, go to [Step 5](#).

Step 3 Adjust the number of tasks submitted to the background thread pool. (For example, cancel some time-consuming tasks with low performance.)

Step 4 Check whether the values of Background Thread Count and Background Thread Usage decrease.

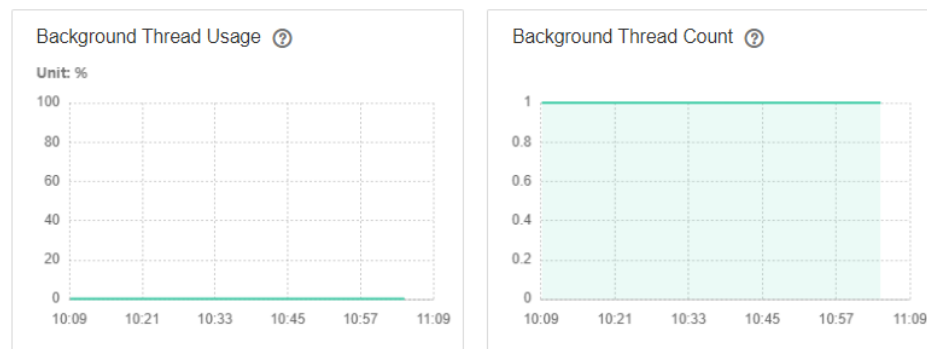
- If it is, go to [Step 7](#).
- If it is not, go to [Step 5](#).

Check the capacity of the HiveServer background thread pool.

Step 5 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive**. On the displayed page, click **HiveServer Instance** and check values of Background Thread Count and Background Thread Usage.

Figure 7-79 Background

Chart



Step 6 Increase the value of `hive.server2.async.exec.threads` in the `${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/1_23_HiveServer/etc/hive-site.xml` file. For example, increase the value by 20%.

Step 7 Save the modification.


Step 8 Check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 9](#).

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 10 Select **Hive** in the required cluster from the **Service**.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.126 ALM-16004 Hive Service Unavailable

Description

This alarm is generated when the HiveServer service is unavailable. The system checks the HiveServer service status every 60 seconds.

This alarm is cleared when the HiveServer service is normal.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 16004 | Critical | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

The system cannot provide data loading, query, and extraction services.

Possible Causes

- Hive service unavailability may be related to the faults of the Hive process as well as basic services, such as ZooKeeper, Hadoop distributed file system (HDFS), Yarn, and DBService.
 - The ZooKeeper service is abnormal.
 - The HDFS service is abnormal.

- The Yarn service is abnormal.
- The DBService service is abnormal.
- The Hive service process is abnormal. If the alarm is caused by Hive process fault, the alarm report has a delay of about 5 minutes.
- The network communication between the Hive and basic services is interrupted.
- The permission on the HDFS temporary directory of Hive is abnormal.
- The local disk space of the Hive node is insufficient.

Procedure

Check the HiveServer/MetaStore process status.

- Step 1** On the FusionInsight Manager portal, click **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance**. In the Hive instance list, check whether the HiveServer or MetaStore instances are in the Unknown state.
- If yes, go to [Step 2](#).
 - If no, go to [Step 4](#).
- Step 2** In the Hive instance list, choose **More** > **Restart Instance** to restart the HiveServer/MetaStore process.

NOTICE

During HiveServer or MetaStore instance restart, the instance cannot provide services for external systems. SQL tasks that are being executed on the instance may fail.

- Step 3** In the alarm list, check whether **Hive Service Unavailable** is cleared.
- If yes, no further action is required.
 - If no, go to [Step 4](#).

Check the ZooKeeper service status.

- Step 4** On the FusionInsight Manager, check whether the alarm list contains **Process Fault**.
- If yes, go to [Step 5](#).
 - If no, go to [Step 8](#).
- Step 5** In the **Process Fault**, check whether **ServiceName** is **ZooKeeper**.
- If yes, go to [Step 6](#).
 - If no, go to [Step 8](#).
- Step 6** Rectify the fault by following the steps provided in "ALM-12007 Process Fault".
- Step 7** In the alarm list, check whether **Hive Service Unavailable** is cleared.
- If yes, no further action is required.
 - If no, go to [Step 8](#).

Check the HDFS service status.

Step 8 On the FusionInsight Manager, check whether the alarm list contains **HDFS Service Unavailable**.

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

Step 9 Rectify the fault by following the steps provided in "ALM-14000 HDFS Service Unavailable".

Step 10 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

Check the Yarn service status.

Step 11 In FusionInsight Manager alarm list, check whether **Yarn Service Unavailable** is generated.

- If yes, go to [Step 12](#).
- If no, go to [Step 14](#).

Step 12 Rectify the fault. For details, see "ALM-18000 Yarn Service Unavailable".

Step 13 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 14](#).

Check the DBService service status.

Step 14 In FusionInsight Manager alarm list, check whether **DBService Service Unavailable** is generated.

- If yes, go to [Step 15](#).
- If no, go to [Step 17](#).

Step 15 Rectify the fault. For details, see "ALM-27001 DBService Service Unavailable".

Step 16 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 17](#).

Check the network connection between the Hive and ZooKeeper, HDFS, Yarn, and DBService.

Step 17 On the FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Hive**.

Step 18 Click **Instance**.

The HiveServer instance list is displayed.

Step 19 Click **Host Name** in the row of **HiveServer**.

The active HiveServer host status page is displayed.

Step 20 Record the IP address under **Basic Information**.

Step 21 Use the IP address obtained in [Step 20](#) to log in to the host where the active HiveServer runs as user **omm**.

Step 22 Run the **ping** command to check whether communication between the host that runs the active HiveServer and the hosts that run the ZooKeeper, HDFS, Yarn, and DBService services is normal. (Obtain the IP addresses of the hosts that run the ZooKeeper, HDFS, Yarn, and DBService services in the same way as that for obtaining the IP address of the active HiveServer.)

- If yes, go to [Step 31](#).
- If no, go to [Step 23](#).

Step 23 Contact the administrator to restore the network.

Step 24 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 31](#).

Check the permission on the HDFS temporary directory.

Step 25 Log in to the node where the HDFS client is located and run the following command to go to the HDFS client installation directory:

```
cd Client installation directory
```

```
source bigdata_env
```

```
kinit user with the supergroup permission (Skip this step for common clusters.)
```

Step 26 Run the following command to check whether the permission on the data warehouse directory is 770:

```
hdfs dfs -ls /tmp | grep hive-scratch
```

- If yes, go to [Step 29](#).
- If no, go to [Step 27](#).

Step 27 Run the following command to restore the default data warehouse permission:

```
hdfs dfs -chmod 770 /tmp/hive-scratch
```

Step 28 Wait for several minutes and check whether the Hive Service Unavailable alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 29](#).

Check whether the local disk space is normal.


Step 29 Run the **df -h** command to check the root directory and check whether the disk usage of the **/srv**, **/var**, and **/opt** directories exceeds 95%.

- If yes, go to [Step 30](#).
- If no, go to [Step 31](#).

Step 30 Clear unnecessary information in the corresponding directory to ensure that the available disk space is greater than 80%. Wait for several minutes and check whether the Hive Service Unavailable alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 31](#).

Collect fault information.

- Step 31** On the FusionInsight Manager, choose **O&M > Log > Download**.
- Step 32** Select the following nodes in the required cluster from the **Service**:
- ZooKeeper
 - HDFS
 - Yarn
 - DBService
 - Hive
- Step 33** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 34** Contact the O&M personnel and send the collected logs.
- End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.127 ALM-16005 The Heap Memory Usage of the Hive Process Exceeds the Threshold

Description

The system checks the Hive service status every 30 seconds. The alarm is generated when the heap memory usage of an Hive service exceeds the threshold (95% of the maximum memory).

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > Hive** to change the threshold.

The alarm is cleared when the heap memory usage is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 16005 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |

Impact on the System

When the heap memory usage of Hive is overhigh, the performance of Hive task operation is affected. In addition, a memory overflow may occur so that the Hive service is unavailable.

Possible Causes

The heap memory of the Hive instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check heap memory usage.

- Step 1** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **Alarm ID** is **16005**. Then check the role name in **Location** and confirm the IP address of the instance.
 - If the role for which the alarm is generated is HiveServer, go to [Step 2](#).
 - If the role for which the alarm is generated is MetaStore, go to [Step 3](#).
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Hive > Instance** and click the HiveServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize > CPU and Memory**, and select **HiveServer Memory Usage Statistics** and click **OK**, check whether the used heap memory of the HiveServer service reaches the threshold (default value: 95%) of the maximum heap memory specified for HiveServer.
 - If yes, go to [Step 4](#).
 - If no, go to [Step 7](#).

Figure 7-80 HiveServer Memory Usage Statistics

Customize Statistics


You have selected 8 items. 2 more items can be selected.

| | | |
|---------------------------|---|--|
| Background (2) | <input checked="" type="checkbox"/> HiveServer Memory Usage Statistics | <input type="checkbox"/> Statistics for the CPU used by the MapReduce jobs submitted by HiveServer |
| CPU and Memory (1) | <input type="checkbox"/> Statistics for the memory used by the MapReduce jobs submitted by HiveServer | <input type="checkbox"/> Direct Memory of HiveServer |
| GC | <input type="checkbox"/> Statistics for the CPUs Used by the HiveServer Process | <input type="checkbox"/> Statistics for the Memory Used by the HiveServer Process |
| HQL (2) | | |
| Session (2) | | |
| Other (1) | | |

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the MetaStore for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **MetaStore Memory Usage Statistics** and click **OK**, check whether the used heap memory of the MetaStore service reaches the threshold (default value: 95%) of the maximum heap memory specified for MetaStore.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 7-81 MetaStore Memory Usage Statistics



Step 4 On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**. Choose **HiveServer/MetaStore > JVM**. Adjust the value of **-Xmx** in **HIVE_GC_OPTS/METASTORE_GC_OPTS** as the following rules. Click **Save**.

NOTE

Suggestions for GC parameter settings for the HiveServer:

- When the heap memory used by the HiveServer process reaches the threshold (default value: 95%) of the maximum heap memory set by the HiveServer process, change the value of **-Xmx** to twice the default value. For example, if **-Xmx** is set to 2GB by default, change the value of **-Xmx** to 4GB. You are advised to change the value of **-Xms** to set the ratio of **-Xms** and **-Xmx** to 1:2 to avoid performance problems when JVM dynamically. On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Hive > CPU and Memory > HiveServer Heap Memory Usage Statistics (HiveServer)** to view **Threshold**.

Suggestions for GC parameter settings for the MetaServer:

- When the heap memory used by the MetaStore process reaches the threshold (default value: 95%) of the maximum heap memory set by the MetaStore process, change the value of **-Xmx** to twice the default value. For example, if **-Xmx** is set to 2GB by default, change the value of **-Xmx** to 4GB. On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Hive > CPU and Memory > MetaStore Heap Memory Usage Statistics (MetaStore)** to view **Threshold**.
- You are advised to change the value of **-Xms** to set the ratio of **-Xms** and **-Xmx** to 1:2 to avoid performance problems when JVM dynamically.

Step 5 Click **More > Restart Service** to restart the service.

NOTICE

During Hive service restart, instances cannot provide services for external systems, and the SQL tasks that are being executed on the instances may fail.


Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 8 Select **Hive** in the required cluster from the **Service**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.128 ALM-16006 The Direct Memory Usage of the Hive Process Exceeds the Threshold

Description

The system checks the Hive service status every 30 seconds. The alarm is generated when the direct memory usage of an Hive service exceeds the threshold (95% of the maximum memory).

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > Hive** to change the threshold.

The alarm is cleared when the direct memory usage is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 16006 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

When the direct memory usage of Hive is overhigh, the performance of Hive task operation is affected. In addition, a memory overflow may occur so that the Hive service is unavailable.

Possible Causes


The direct memory of the Hive instance on the node is overused or the direct memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check direct memory usage.

- Step 1** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **Alarm ID** is **16006**. Then check the role name in **Location** and confirm the IP address of the instance.
- If the role for which the alarm is generated is HiveServer, go to [Step 2](#).
 - If the role for which the alarm is generated is MetaStore, go to [Step 3](#).
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Hive > Instance** and click the HiveServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize > CPU and Memory**, and select **HiveServer Memory Usage Statistics** and click **OK**, check whether the used direct memory of the HiveServer service reaches the threshold (default value: 95%) of the maximum direct memory specified for HiveServer.
- If yes, go to [Step 4](#).
 - If no, go to [Step 7](#).

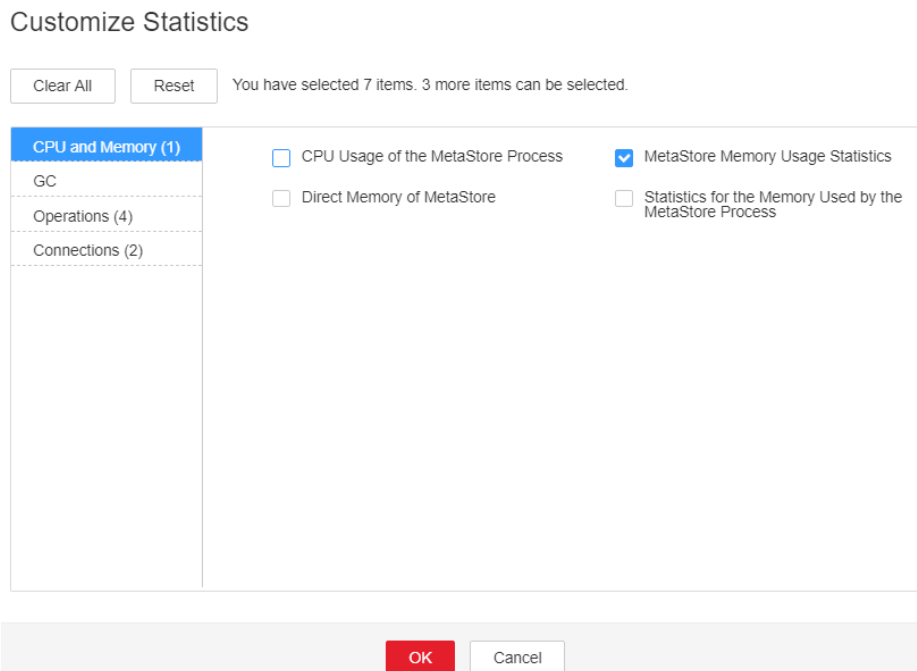
Figure 7-82 HiveServer Memory Usage Statistics



Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the MetaStore for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **MetaStore Memory Usage Statistics** and click **OK**, check whether the used direct memory of the MetaStore service reaches the threshold(default value: 95%) of the maximum direct memory specified for MetaStore.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 7-83 MetaStore Memory Usage Statistics



Step 4 On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**. Choose **HiveServer/MetaStore > JVM**. Adjust the value of **-XX:MaxDirectMemorySize** in **HIVE_GC_OPTS/METASTORE_GC_OPTS** as the following rules. Click **Save**.

NOTE

Suggestions for GC parameter settings for the HiveServer:

- It is recommended that you set the value of **-XX:MaxDirectMemorySize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 8 GB, **-XX:MaxDirectMemorySize** is set to 1024 MB. If **-Xmx** is set to 4 GB, **-XX:MaxDirectMemorySize** is set to 512 MB. It is recommended that the value of **-XX:MaxDirectMemorySize** be greater than or equal to 512 MB.

Suggestions for GC parameter settings for the MetaServer:

- It is recommended that you set the value of **-XX:MaxDirectMemorySize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 8 GB, **-XX:MaxDirectMemorySize** is set to 1024 MB. If **-Xmx** is set to 4 GB, **-XX:MaxDirectMemorySize** is set to 512 MB. It is recommended that the value of **-XX:MaxDirectMemorySize** be greater than or equal to 512 MB.

Step 5 Click **More > Restart Service** to restart the service.


NOTICE

During Hive service restart, instances cannot provide services for external systems, and the SQL tasks that are being executed on the instances may fail.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

- Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 8** Select **Hive** in the required cluster from the **Service**.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.129 ALM-16007 Hive GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) time of the Hive service every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (exceeds 12 seconds for three consecutive checks.) To change the threshold, choose **O&M > Alarm > Thresholds > Name of the desired cluster > Hive**. This alarm is cleared when the Hive GC time is shorter than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 16007 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |

| Name | Meaning |
|-------------------|--|
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

If the GC time exceeds the threshold, Hive data read and write are affected, task execution may slow down, or services may restart unexpectedly.

Possible Causes


The memory of Hive instances is overused, the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC time.

- Step 1** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **Alarm ID** is **16007**. Then check the role name in **Location** and confirm the IP address of the instance.
- If the role for which the alarm is generated is HiveServer, go to [Step 2](#).
 - If the role for which the alarm is generated is MetaStore, go to [Step 3](#).
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Hive > Instance** and click the HiveServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize > GC**, and select **Garbage Collection (GC) Time of HiveServer** and click **OK** to check whether the GC time is longer than 12 seconds.
- If yes, go to [Step 4](#).
 - If no, go to [Step 7](#).

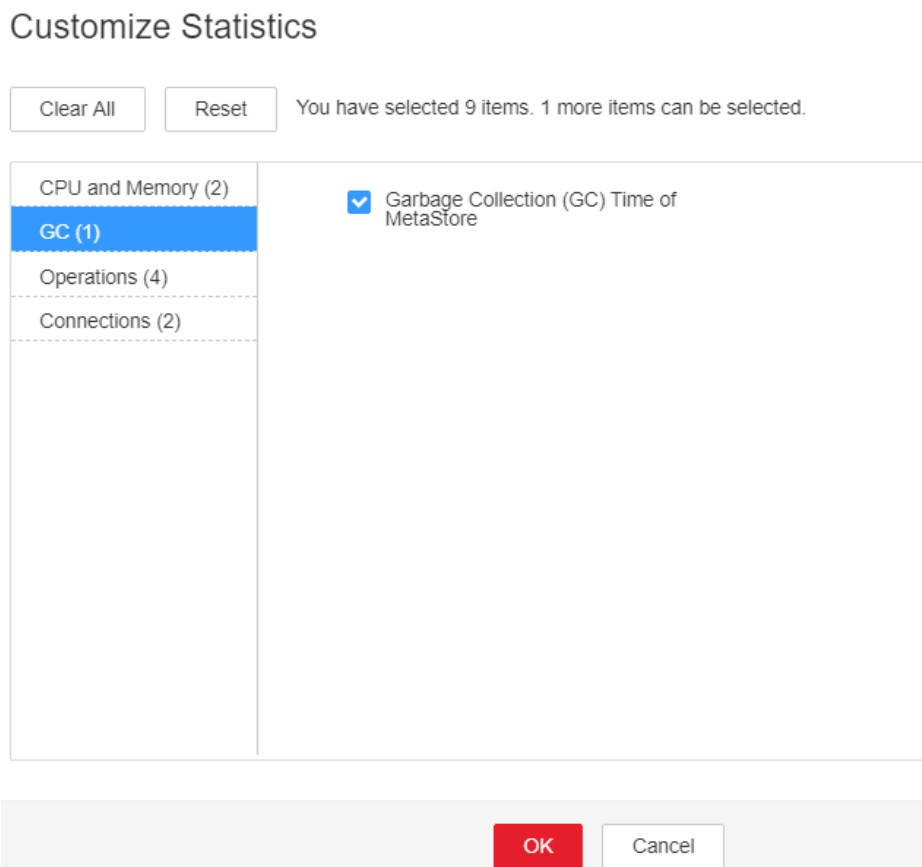
Figure 7-84 Garbage Collection (GC) Time of HiveServer
Customize Statistics



Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the MetaStore for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **GC**, and select **Garbage Collection (GC) Time of MetaStore** and click **OK** to check whether the GC time is longer than 12 seconds.

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

Figure 7-85 Garbage Collection (GC) Time of MetaStore



Check the current JVM configuration.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations**. Choose **HiveServer/MetaStore** > **JVM**. Adjust the value of **-Xmx** in **HIVE_GC_OPTS/METASTORE_GC_OPTS** as the following rules. Click **Save**.

NOTE

Suggestions for GC parameter settings for the HiveServer:

- When the Hive GC time exceeds the threshold, change the value of **-Xmx** to twice the default value. For example, if **-Xmx** is set to 2 GB by default, change the value of **-Xmx** to 4 GB.
- You are advised to change the value of **-Xms** to set the ratio of **-Xms** and **-Xmx** to 1:2 to avoid performance problems when JVM dynamically.

Suggestions for GC parameter settings for the MetaServer:

- When the Meta GC time exceeds the threshold, change the value of **-Xmx** to twice the default value. For example, if **-Xmx** is set to 2 GB by default, change the value of **-Xmx** to 4 GB.
- You are advised to change the value of **-Xms** to set the ratio of **-Xms** and **-Xmx** to 1:2 to avoid performance problems when JVM dynamically.

Step 5 Click **More** > **Restart Service** to restart the service.

NOTICE

During Hive service restart, instances cannot provide services for external systems, and the SQL tasks that are being executed on the instances may fail.


Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Collect fault information.

Step 7 On the FusionInsight Manager portal of active and standby clusters, choose **O&M > Log > Download**.

Step 8 In the **Service**, select **Hive** in the required cluster.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.130 ALM-16008 Non-Heap Memory Usage of the Hive Process Exceeds the Threshold

Description

The system checks the Hive service status every 30 seconds. The alarm is generated when the non-heap memory usage of an Hive service exceeds the threshold (95% of the maximum memory).

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > Hive** to change the threshold.

The alarm is cleared when the non-heap memory usage is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 16008 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |

Impact on the System

When the non-heap memory usage of Hive is overhigh, the performance of Hive task operation is affected. In addition, a memory overflow may occur so that the Hive service is unavailable.

Possible Causes

The non-heap memory of the Hive instance on the node is overused or the non-heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check non-heap memory usage.

- Step 1** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **Alarm ID** is **16008**. Then check the role name in **Location** and confirm the IP address of the instance.
 - If the role for which the alarm is generated is HiveServer, go to [Step 2](#).
 - If the role for which the alarm is generated is MetaStore, go to [Step 3](#).
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Hive > Instance** and click the HiveServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize > CPU and Memory**, and select **HiveServer Memory Usage Statistics** and click **OK**, check whether the used non-heap memory of the HiveServer service reaches the threshold(default value: 95%) of the maximum non-heap memory specified for HiveServer.
 - If yes, go to [Step 4](#).
 - If no, go to [Step 7](#).

Figure 7-86 HiveServer Memory Usage Statistics


Customize Statistics

You have selected 8 items. 2 more items can be selected.

| | | |
|---------------------------|---|--|
| Background (2) | <input checked="" type="checkbox"/> HiveServer Memory Usage Statistics | <input type="checkbox"/> Statistics for the CPU used by the MapReduce jobs submitted by HiveServer |
| CPU and Memory (1) | <input type="checkbox"/> Statistics for the memory used by the MapReduce jobs submitted by HiveServer | <input type="checkbox"/> Direct Memory of HiveServer |
| GC | <input type="checkbox"/> Statistics for the CPUs Used by the HiveServer Process | <input type="checkbox"/> Statistics for the Memory Used by the HiveServer Process |
| HQL (2) | | |
| Session (2) | | |
| Other (1) | | |

- Step 3** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the MetaStore for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **MetaStore Memory Usage Statistics** and click **OK**, check whether the used non-heap memory of the MetaStore service reaches the threshold (default value: 95%) of the maximum non-heap memory specified for MetaStore.
- If yes, go to [Step 4](#).
 - If no, go to [Step 7](#).

Figure 7-87 MetaStore Memory Usage Statistics



Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations**. Choose **HiveServer/MetaStore** > **JVM**. Adjust the value of **-XX:MaxMetaspaceSize** in **HIVE_GC_OPTS/METASTORE_GC_OPTS** as the following rules. Click **Save**.

NOTE

Suggestions for GC parameter settings for the HiveServer:

- It is recommended that you set the value of **-XX:MaxMetaspaceSize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 2 GB, **-XX:MaxMetaspaceSize** is set to 256 MB. If **-Xmx** is set to 4 GB, **-XX:MaxMetaspaceSize** is set to 512 MB.

Suggestions for GC parameter settings for the MetaServer:

- It is recommended that you set the value of **-XX:MaxMetaspaceSize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 2 GB, **-XX:MaxMetaspaceSize** is set to 256 MB. If **-Xmx** is set to 4 GB, **-XX:MaxMetaspaceSize** is set to 512 MB.

Step 5 Click **More** > **Restart Service** to restart the service.

NOTICE

During Hive service restart, instances cannot provide services for external systems, and the SQL tasks that are being executed on the instances may fail.


Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 8 Select **Hive** in the required cluster from the **Service**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.131 ALM-16009 Map Number Exceeds the Threshold

Description

The system checks the number of HQL maps in every 30 seconds. This alarm is generated if the number exceeds the threshold. By default, **Trigger Count** is set to **3**, and the threshold is 5000.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16009 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

| Name | Meaning |
|-------------------|--|
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

If the number of HQL maps executed by Hive is too large, a large number of Yarn queue resources are occupied, which may take a long time and affect other tasks running using this queue.

Possible Causes


The HQL statements are not the optimal.

Procedure

Check the number of HQL maps.

- Step 1** On FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Resource**. Check the HQL statements with the excessively large number (5000 or more) of maps in **HQL Map Count**.
- Step 2** Locate the corresponding HQL statements, optimize them and execute them again.
- Step 3** Check whether the alarm is cleared.
 - If it is, no further action is required.
 - If it is not, go to [Step 4](#).

Collect fault information.

- Step 4** On the FusionInsight Manager, choose **O&M** > **Log** > **Download**.
- Step 5** Select **Hive** in the required cluster from the **Service**.
- Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 7** Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.132 ALM-16045 Hive Data Warehouse Is Deleted

Description

The system checks the Hive data warehouse in every 60 seconds. This alarm is generated when the Hive data warehouse is deleted.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16045 | Critical | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

The default Hive data warehouse is deleted. As a result, creating databases or tables in the default data warehouse fails, and services are affected.

Possible Causes

Hive periodically checks the status of the default data warehouse and finds that the default data warehouse is deleted.

Procedure

Check the default Hive data warehouse.

Step 1 Log in to the node where the client is located as user **root**.

Step 2 Run the following command to check whether the **warehouse** directory exists in **hdfs://hacluster/user/<username>.Trash/Current/**.

```
hdfs dfs -ls hdfs://hacluster/user/<username>.Trash/Current/
```

For example, if **user/hive/warehouse** exists:

```
host01:/opt/client # hdfs dfs -ls hdfs://hacluster/user/test/.Trash/Current/  
Found 1 items  
drwx----- - test hadoop 0 2019-06-17 19:53 hdfs://hacluster/user/test/.Trash/Current/user
```

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 By default, there is an automatic recovery mechanism for the data warehouse. You can wait for 5 ~10s to check whether the default data warehouse is restored. If the data warehouse is not recovered, manually run the following command to restore the data warehouse.

```
hdfs dfs -mv hdfs://hacluster/user/<username>/.Trash/Current/user/hive/  
warehouse /user/hive/warehouse
```

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 Collect related information in the **.Trash/Current/** directory on the client background.

Step 6 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.133 ALM-16046 Hive Data Warehouse Permission Is Modified

Description

The system checks the Hive data warehouse permission in every 60 seconds. This alarm is generated if the permission is modified.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16046 | Critical | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

If the permission on the Hive default data warehouse is modified, the permission for users or user groups to create databases or tables in the default data warehouse is changed.

Possible Causes

Hive periodically checks the status of the default data warehouse and finds that default data warehouse permission is changed.

Procedure

Check the Hive default data warehouse permission.

Step 1 Log in to the node where the client is located as user **root**.

Step 2 Run the following command to go to the HDFS client installation directory:

```
cd Client installation directory
```

```
source bigdata_env
```

```
kinit User who has the supergroup permission (Skip this step for a common cluster.)
```

Step 3 Run the following command to restore the default data warehouse permission:

- Security mode: **hdfs dfs -chmod 770 hdfs://hacluster/user/hive/warehouse**
- Non-security mode: **hdfs dfs -chmod 777 hdfs://hacluster/user/hive/warehouse**

Step 4 Check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 5](#).

Collect fault information.

Step 5 Collect related information in the **hdfs://hacluster/user/hive/warehouse** directory on the client background.

Step 6 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.134 ALM-16047 HiveServer Has Been Deregistered from ZooKeeper

Alarm Description

The system checks the Hive service every 60 seconds. This alarm is generated when Hive registration information on ZooKeeper is lost or Hive cannot connect to ZooKeeper.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 16047 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |

Impact on the System

When a Hive client sets up a new connection, it cannot select the HiveServer node that has been deregistered from ZooKeeper. If all HiveServer nodes have been deregistered from ZooKeeper, the HiveServer service will be unavailable.

Possible Causes

- The ZooKeeper instance is abnormal.
- Some Hive configurations are incorrect.

Handling Procedure

Check the ZooKeeper service status.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms** and check whether **ALM-12007 Process Fault** exists in the alarm list.

- If yes, go to [Step 2](#).
- If no, go to [Step 5](#).

Step 2 In **Location** of **ALM-12007 Process Fault**, check whether the service name is **ZooKeeper**.


- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Rectify the fault by following steps provided in **ALM-12007 Process Fault**.

Step 4 In the alarm list, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check whether the Hive configurations are correctly modified.

Step 5 On FusionInsight Manager, choose **Audit**. On the **Audit** page, click **Advanced Search**, click  on the right of **Operation Type**, select **Save configuration**, click **OK**, and click **Search**.

Step 6 In the search result, check the historical configurations of Hive- and ZooKeeper-related services in the **Service** column. [Table 7-103](#) lists some configurations that may affect the connection between Hive and ZooKeeper.

Table 7-103 Configurations related to connection between Hive and ZooKeeper

| Service | Parameter | Description |
|---------|----------------------------|--|
| Hive | HIVE_GC_OPTS | HiveServer memory configuration. If the configuration is abnormal, HiveServer may restart repeatedly. In this case, you need to check the health status of the instance processes. |
| | hive.zookeeper.quorum | IP address of the node accommodating ZooKeeper that is connected to Hive. |
| | hive.zookeeper.client.port | Port of the ZooKeeper client connected to Hive. |

| Service | Parameter | Description |
|-----------|---------------------------------------|--|
| | hive.zookeeper.session.timeout | Timeout interval of the session set up between Hive and ZooKeeper. |
| | hive.zookeeper.connection.timeout | Timeout interval for Hive to connect to ZooKeeper. |
| | hive.zookeeper.connection.max.retries | Maximum number of retries for Hive to connect to ZooKeeper. |
| ZooKeeper | clientPort | Port number of the ZooKeeper client. |
| | ssl.enabled | Whether to enable SSL connections of ZooKeeper. |

Restart related instances.

Step 7 Log in to FusionInsight Manager. Choose **O&M > Alarm > Alarms**, click the drop-down list in the row that contains the alarm, and view the role and the IP address of the node for which the alarm is generated in **Location**.

Step 8 Choose **Cluster**, click the name of the desired cluster, and choose **Services > Hive > Instance**. On the page that is displayed, select the instance at the IP address for which the alarm is generated, click **More**, and select **Restart Instance**.

NOTICE

During Hive instance restart, the instance cannot provide services for external systems. SQL tasks that are being executed on the instance may fail.


Step 9 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 11 Expand the **Service** drop-down list, and select **Hive** for the target cluster.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.135 ALM-16048 Tez or Spark Library Path Does Not Exist

Description

The system checks the Tez and Spark library paths every 180 seconds. This alarm is generated when the Tez or Spark library path does not exist.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16048 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

The Hive on Tez and Hive on Spark functions are affected.

Possible Causes

The Tez or Spark library path is deleted from the HDFS.

Procedure

Check the default Hive data warehouse.

Step 1 Log in to the node where the client is located as user **root**.

Step 2 Run the following command to check whether the **tezlib** or **sparklib** directory exists in the **hdfs://hacluster/user/{User name}/.Trash/Current/** director:

```
hdfs dfs -ls hdfs://hacluster/user/<username>/.Trash/Current/
```

For example, the following information shows that `/user/hive/tezlib/8.1.0.1/` and `/user/hive/sparklib/8.1.0.1/` exist.

```
host01:/opt/client # hdfs dfs -ls hdfs://hacluster/user/test/.Trash/Current/
Found 1 items
drwx----- - test hadoop 0 2019-06-17 19:53 hdfs://hacluster/user/test/.Trash/Current/user
```

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Run the following command to restore **tezlib** and **sparklib**.

```
hdfs dfs -mv hdfs://hacluster/user/<username>/.Trash/Current/user/hive/tezlib/8.1.0.1/tez.tar.gz /user/hive/tezlib/8.1.0.1/tez.tar.gz
```

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 Collect related information in the `.Trash/Current/` directory on the client background.

Step 6 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.136 ALM-17003 Oozie Service Unavailable

Description

The system checks the Oozie service status in every 5 seconds. This alarm is generated when Oozie or a component on which Oozie depends cannot provide services properly.

This alarm is automatically cleared when the Oozie service recovers.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 17003 | Critical | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

Oozie cannot be used to schedule jobs.

Possible Causes

- The DBService service is abnormal or the data of Oozie stored in DBService is damaged.
- The HDFS service is abnormal or the data of Oozie stored in HDFS is damaged.
- The Yarn service is abnormal.
- The Nodeagent process is abnormal.

Procedure

Query the Oozie service health status code.

- Step 1** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Oozie**. Click **oozie** (any one is OK) on the **oozie WebUI**. to go to the Oozie WebUI.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

- Step 2** Add **/servicehealth** to the URL in the address box of the browser and access again. The value of **statusCode** is the current Oozie service health status code.

For example, visit **https://10.10.0.117:20026/Oozie/oozie/130/oozie/servicehealth**. The result is as follows:

```
{"beans":[{"name":"serviceStatus","statusCode":0}]}
```

If the health status code cannot be displayed or the browser does not respond, the service may be unavailable due to Oozie process fault. See **Step 13** to rectify the fault.

Step 3 Perform the operations based on the error code. For details, see [Table 7-104](#).

Table 7-104 Oozie service health status code

| Status Code | Description | Error Cause | Solution |
|-------------|------------------------------------|---|-------------------------------|
| 0 | The service is running properly. | None | None |
| 18002 | The DBService service is abnormal. | Oozie fails to connect to DBService or the data stored in DBService is damaged. | See Step 4 . |
| 18003 | The HDFS service is abnormal. | Oozie fails to connect to HDFS or the data stored in HDFS is damaged. | See Step 7 . |
| 18005 | The MapReduce service is abnormal. | The Yarn service is abnormal. | See Step 11 . |

Check the DBService service.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the DBService service is running properly.

- If yes, go to [Step 6](#).
- If no, go to [Step 5](#).

Step 5 Resolve the problem of DBService based on the alarm help and check whether the Oozie alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 18](#).

Step 6 Log in to the Oozie database to check whether the data is complete.

1. Log in to the active DBService node as user **root**.

On the FusionInsight Manager page, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService** > **Instance** to view the IP address of the active DBService node.

2. Run the following command to log in to the Oozie database:

```
su - omm
```

```
source ${BIGDATA_HOME}/FusionInsight_BASE_8.1.0.1/install/  
FusionInsight-dbservice-2.7.0/dbservice_profile
```

```
gsql -U Username -W Oozie database password -p 20051 -d Database name
```

3. After the login is successful, enter **\d** to check whether there are 15 data tables.

The Oozie service has 15 data tables by default. If these data tables are deleted or the table structure is modified, the Oozie service may be unavailable. Contact the O&M personnel to back up the data and perform restoration.

Check the HDFS service.

Step 7 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the HDFS service is running properly.

- If yes, go to [Step 9](#).
- If no, go to [Step 8](#).

Step 8 Resolve the problem of HDFS based on the alarm help and check whether the Oozie alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 18](#).

Step 9 Log in to HDFS to check whether the Oozie file directory structure is complete.

1. Download and install an HDFS client..
2. Log in to the client node as user **root** and run the following commands to check whether **/user/oozie/share** exists.

If the cluster uses the security mode, perform security authentication.

```
kinit admin
```

```
hdfs dfs -ls /user/oozie/share
```

- If yes, go to [Step 18](#).
- If no, go to [Step 10](#).

Step 10 In the Oozie client installation directory, manually upload the share directory to **/user/oozie** in HDFS, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 18](#).

Check the Yarn and MapReduce service.

Step 11 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the Yarn and MapReduce services are running properly.

- If yes, go to [Step 18](#).
- If no, go to [Step 12](#).

Step 12 Resolve the problem of Yarn and MapReduce based on the alarm help and check whether the Oozie alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 18](#).

Check the Oozie process.

Step 13 Log in to each node of Oozie as user **root**.

Step 14 Run the **ps -ef | grep oozie** command to check whether the Oozie process exists.

- If yes, go to [Step 15](#).
- If no, go to [Step 18](#).

Step 15 Collect fault information in **prestartDetail.log**, **oozie.log**, and **catalina.out** in the Oozie log directory **/var/log/Bigdata/oozie**. If the alarm is not caused by manual misoperation, go to [Step 16](#).

Check the Nodeagent process.

Step 16 Log in to each node of Oozie as user **root**. Run the **ps -ef | grep nodeagent** command to check whether the Nodeagent process exists.

- If yes, go to [Step 17](#).
- If no, go to [Step 18](#).

Step 17 Run the **kill -9 *The process ID of nodeagent*** command, wait 10 minutes, and check whether alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 18](#).

Step 18 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.137 ALM-17004 Oozie Heap Memory Usage Exceeds the Threshold

Description

The system checks the heap memory usage of the Oozie service every 60 seconds. The alarm is generated when the heap memory usage of a Metadata instance exceeds the threshold (95% of the maximum memory). The alarm is cleared when the heap memory usage is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 17004 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

The heap memory overflow may cause a service breakdown. After the service breaks down, the Oozie service cannot be used to schedule tasks.

Possible Causes

The heap memory of the Oozie instance is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Oozie Heap Memory Usage Exceeds the Threshold > Location**. Check the IP address of the instance involved in this alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Oozie > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize > Memory > Oozie Heap Memory Resource Percentage**. Click **OK**.

Figure 7-88 Oozie Heap Memory Resource Percentage

Customize Statistics

You have selected 4 items. 5 more items can be selected.

| | | |
|------------|---|---|
| GC (3) | <input checked="" type="checkbox"/> Oozie Heap Memory Resource Percentage | <input type="checkbox"/> Oozie Heap Memory Resource |
| Memory (1) | <input type="checkbox"/> Oozie Direct Buffer Resource Percentage | <input type="checkbox"/> Oozie Direct Buffer Resource |
| | <input type="checkbox"/> Oozie Non Heap Memory Resource Percentage | <input type="checkbox"/> Oozie Non Heap Memory Resource |

Step 3 Check whether the used heap memory of Oozie reaches the threshold (the default value is 95% of the maximum heap memory) specified for Oozie.

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Oozie** > **Configurations** > **All Configurations**. Set Search **GC_OPTS** in the search box. Increase the value of **-Xmx** as required, and click **Save** > **OK**.

 **NOTE**

Suggestions on GC parameter settings for Oozie:

You are advised to set **-Xms** and **-Xmx** to the same value to prevent adverse impact on performance when JVM dynamically adjusts the heap memory size.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

NOTICE

During the service or instance restart, services are interrupted, but submitted jobs are not affected.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

- Step 7** Select **Oozie** in the required cluster from the **Service**.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected logs.
- End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.138 ALM-17005 Oozie Non Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the non heap memory usage of Oozie every 30 seconds. This alarm is reported if the non heap memory usage of Oozie exceeds the threshold (80%). This alarm is cleared if the non heap memory usage is lower than the threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 17005 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

The service may break down and the Oozie service cannot be used to schedule tasks.

Possible Causes

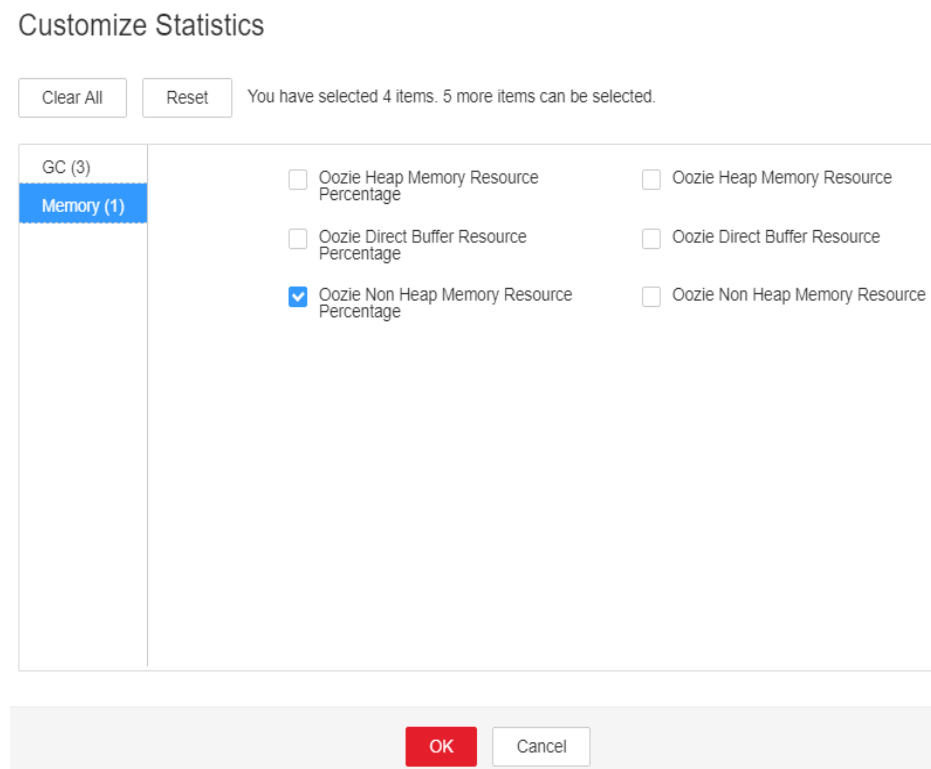
The non-heap memory of the Oozie instance is overused or the non-heap memory is inappropriately allocated.

Handling Procedure

Check non-heap memory usage.

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > Oozie Non Heap Memory Usage Exceeds the Threshold**. On the displayed page, check the location information of the alarm. Check the name of the instance host for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Name of the target cluster > Services > Oozie** and click the **Instance** tab. On the displayed page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Memory** and select **Oozie Non Heap Memory Resource Percentage**. Click **OK**.

Figure 7-89 Oozie non-heap memory usage



Step 3 Check whether the non-heap memory used by Oozie reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the target cluster* > **Services** > **Oozie** and click the **Configurations** and then **All Configurations**. On the displayed page, search for the **GC_OPTS** parameter in the search box and check whether it contains **-XX: MaxMetaspaceSize**. If yes, increase the value of **-XX: MaxMetaspaceSize** based on the site requirements. If no, manually add **-XX: MaxMetaspaceSize** and set its value to 1/8 of the value of **-Xmx**. Click **Save**, and then click **OK**.

 **NOTE**

JDK1.8 does not support the **MaxPermSize** parameter.

Suggestions on GC parameter settings for Oozie:

Set the value of **-XX:MaxMetaspaceSize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 2 GB, **-XX:MaxMetaspaceSize** is set to 256 MB. If **-Xmx** is set to 4 GB, **-XX:MaxMetaspaceSize** is set to 512 MB.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).


NOTICE

During the service or instance restart, services are interrupted, but submitted jobs are not affected.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Oozie** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.139 ALM-17006 Oozie Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the Oozie service every 30 seconds. The alarm is generated when the direct memory usage of an Oozie instance exceeds the threshold (80% of the maximum memory). The alarm is cleared when the direct memory usage of Oozie is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 17006 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

The direct memory overflow may cause a service breakdown. After the service breaks down, the Oozie service cannot be used to schedule tasks.

Possible Causes

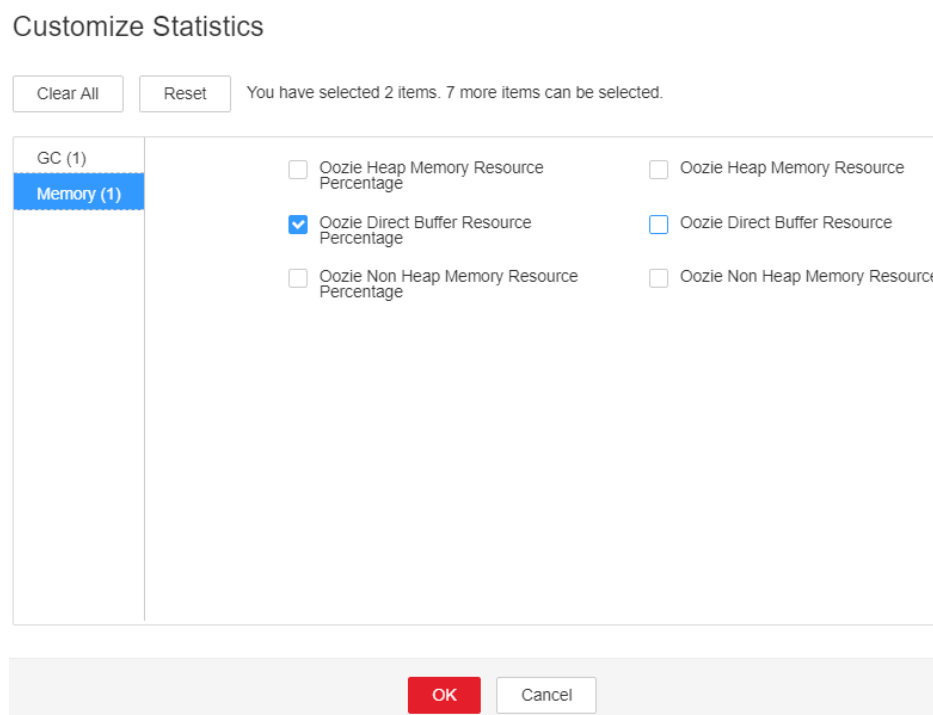
The direct memory of the Oozie instance is overused or the direct memory is inappropriately allocated.

Procedure

Check direct memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Oozie Direct Memory Usage Exceeds the Threshold > Location**. Check the IP address of the instance involved in this alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Oozie > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize > Memory > Oozie Direct Buffer Resource Percentage**. Click **OK**.

Figure 7-90 Oozie Direct Buffer Resource Percentage



- Step 3** Check whether the used direct memory of Oozie reaches the threshold (the default value is 80% of the maximum direct memory) specified for Oozie.
- If yes, go to **Step 4**.
 - If no, go to **Step 6**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Oozie > Configurations**. Click **All Configurations**. Search **GC_OPTS** in the search box. Increase the value of **-XX:MaxDirectMemorySize** as required, and click **Save**. Click **OK**.

NOTE

Suggestions on GC parameter settings for Oozie:

You are advised to set the value of **-XX:MaxDirectMemorySize** to 1/4 of the value of **-Xmx**. For example, if **-Xmx** is set to 4 GB, **-XX:MaxDirectMemorySize** is set to 1024 MB. If **-Xmx** is set to 2 GB, **-XX:MaxDirectMemorySize** is set to 512 MB. It is recommended that the value of **-XX:MaxDirectMemorySize** be greater than or equal to 512 MB.


- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.

- If no, go to [Step 6](#).

NOTICE

During the service or instance restart, services are interrupted, but submitted jobs are not affected.

Collect fault information.

- Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 7** Select **Oozie** in the required cluster from the **Service** drop-down list.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.140 ALM-17007 Garbage Collection (GC) Time of the Oozie Process Exceeds the Threshold

Description

The system checks GC time of the Oozie process every 60 seconds. The alarm is generated when GC time of the Oozie process exceeds the threshold (default value: **12 seconds**). The alarm is cleared when GC time is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 17007 | Major | Yes |

Parameters

| Name | Meaning |
|--------|---|
| Source | Specifies the cluster for which the alarm is generated. |

| Name | Meaning |
|-------------------|--|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

Oozie scheduling task responds slowly until the service is unavailable.

Possible Causes

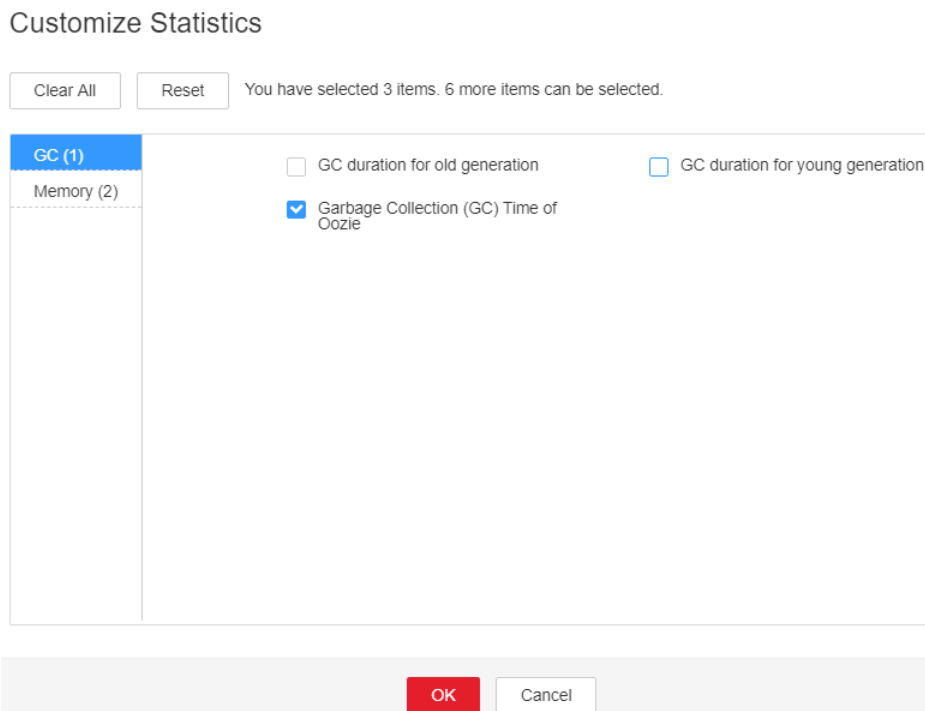
The heap memory of the Oozie instance is overused or the heap memory is inappropriately allocated.

Procedure

Check GC time.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Garbage Collection (GC) Time of the Oozie Process Exceeds the Threshold > Location**. Check the IP address of the instance involved in this alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Oozie > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize > GC > Garbage Collection (GC) Time of Oozie**. Click **OK**.

Figure 7-91 Garbage Collection (GC) Time of Oozie



Step 3 Check whether GC time of the Oozie process every second exceeds the threshold (default value: **12 seconds**).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

Step 4 On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Oozie > Configurations**. Click **All Configurations**. Search **GC_OPTS** in the search box. Increase the value of **-Xmx** as required, and click **Save**. Click **OK**.

NOTE

Suggestions on GC parameter settings for Oozie:

You are advised to set **-Xms** and **-Xmx** to the same value to prevent adverse impact on performance when JVM dynamically adjusts the heap memory size.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

NOTICE

During the service or instance restart, services are interrupted, but submitted jobs are not affected.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

- Step 7** Select **Oozie** in the required cluster from the **Service** drop-down list.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected logs.
- End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.141 ALM-17008 Abnormal Connection Between Oozie and ZooKeeper

Alarm Description

In HA mode, Oozie depends on ZooKeeper. This alarm is generated when the connection between Oozie and ZooKeeper is abnormal for three consecutive times.

This alarm is cleared when the connection between Oozie and ZooKeeper becomes normal.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 17008 | Minor | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

| Parameter | Description |
|-------------------|---|
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

Running scheduling tasks are blocked and new scheduling tasks cannot be submitted. In HA mode, the Oozie service will restart if this alarm is reported.

Possible Causes

- The ZooKeeper service is abnormal.
- Oozie fails to connect to ZooKeeper.

Handling Procedure

Check the ZooKeeper service status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of ZooKeeper is **Normal**.

- If yes, go to [Step 5](#).
- If no, go to [Step 2](#).

Step 2 In the alarm list, check whether **ALM-13000 ZooKeeper Service Unavailable** is reported.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Rectify the fault by performing the operations provided for **ALM-13000 ZooKeeper Service Unavailable**.

Step 4 Wait for several minutes and check whether the alarm **Abnormal Connection Between Oozie and ZooKeeper** is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check the connectivity between Oozie and ZooKeeper.


Step 5 Log in to FusionInsight Manager, choose **O&M > Log > Online Search**, select the Oozie service, and search for the keyword **[Oozie Alarm Enhancement] [ZooKeeper]** in the log. View the cause in the log, and rectify the fault. In the alarm list, check whether the alarm **Abnormal Connection Between Oozie and ZooKeeper** is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 7 Select **Oozie** for **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.142 ALM-17009 Abnormal Connection Between Oozie and DBService

Alarm Description

Oozie depends on DBService. After a task is submitted, the system checks DBService connectivity. This alarm is generated when the service fails the check for 10 consecutive times.

This alarm is cleared when the connection between Oozie and DBService becomes normal.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 17009 | Minor | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

| Parameter | Description |
|-------------------|---|
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

Running scheduling tasks are blocked and new scheduling tasks cannot be submitted.

Possible Causes

- The DBService service is abnormal.
- Oozie fails to connect to DBService.

Handling Procedure

Check the DBService status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of DBService is **Normal**.

- If yes, go to [Step 5](#).
- If no, go to [Step 2](#).

Step 2 In the alarm list, check whether **ALM-27001 DBService Service Unavailable** is reported.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Rectify the fault by performing the operations provided for **ALM-27001 DBService Service Unavailable**.

Step 4 Wait for several minutes and check whether the alarm **Abnormal Connection Between Oozie and DBService** is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check the connectivity between Oozie and DBService.


Step 5 Log in to FusionInsight Manager, choose **O&M > Log > Online Search**, select the Oozie service, and search for the keyword **[Oozie Alarm Enhancement][DB Service]** in the log. View the cause in the log, and rectify the fault. In the alarm list, check whether the alarm **Abnormal Connection Between Oozie and DBService** is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 7 Select **Oozie** for **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.143 ALM-17010 Abnormal Connection Between Oozie and HDFS

Alarm Description

Oozie depends on HDFS. After a task is submitted, the system checks HDFS connectivity. This alarm is generated when the service fails the check for 3 consecutive times.

This alarm is cleared when the connection between Oozie and HDFS becomes normal.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 17010 | Minor | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

| Parameter | Description |
|-------------------|---|
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

Running scheduling tasks are blocked and new scheduling tasks cannot be submitted.

Possible Causes

The HDFS service restarts, there is a fault, or the network connectivity is abnormal.

Handling Procedure

Check the HDFS service status.

- Step 1** In the service list on FusionInsight Manager, check whether **Running Status** of HDFS is **Normal**.
- If yes, go to [Step 5](#).
 - If no, go to [Step 2](#).
- Step 2** In the alarm list, check whether the "ALM-14000 HDFS Service Unavailable" alarm is generated.
- If yes, go to [Step 3](#).
 - If no, go to [Step 5](#).
- Step 3** Rectify the fault by performing the operations provided for **ALM-14000 HDFS Service Unavailable**.
- Step 4** Wait for several minutes and check whether the alarm **Abnormal Connection Between Oozie and HDFS** is cleared.
- If yes, no further action is required.
 - If no, go to [Step 5](#).


Check the connectivity between Oozie and HDFS.

- Step 5** Log in to FusionInsight Manager, choose **O&M > Log > Online Search**, select the Oozie service, and search for the keyword **[Oozie Alarm Enhancement][HDFS]** in the log. View the cause in the log, and rectify the fault. In the alarm list, check whether the alarm **Abnormal Connection Between Oozie and HDFS** is cleared.
- If yes, no further action is required.
 - If no, go to [Step 6](#).

Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 7 Select **Oozie** for **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.144 ALM-17011 Abnormal Connection Between Oozie and Yarn

Alarm Description

Oozie depends on Yarn. After a task is submitted, the system checks Yarn connectivity. This alarm is generated when the service fails the check for 5 consecutive times.

This alarm is cleared when the connection between Oozie and Yarn becomes normal.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 17011 | Minor | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

| Parameter | Description |
|-------------------|---|
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

Running scheduling tasks are blocked and new scheduling tasks cannot be submitted.

Possible Causes

- The Yarn service is abnormal.
- The connection between Oozie and Yarn is abnormal.

Handling Procedure

Check the YARN service status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of Yarn is **Normal**.

- If yes, go to [Step 5](#).
- If no, go to [Step 2](#).

Step 2 In the alarm list, check whether **ALM-18000 YARN Service Unavailable** is generated.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Rectify the fault by performing the operations provided for **ALM-18000 Yarn Service Unavailable**.

Step 4 Wait for several minutes and check whether the alarm **Abnormal Connection Between Oozie and Yarn** is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check the connectivity between Oozie and Yarn.


Step 5 Log in to FusionInsight Manager, choose **O&M > Log > Online Search**, select the Oozie service, and search for the keyword **[Oozie Alarm Enhancement][Yarn]** in the log. View the cause in the log, and rectify the fault. In the alarm list, check whether the alarm **Abnormal Connection Between Oozie and Yarn** is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 7 Select **Oozie** for **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.145 ALM-18000 Yarn Service Unavailable

Description

This alarm is generated when the Yarn service is unavailable. The alarm module checks the Yarn service status every 60 seconds.

The alarm is cleared when the Yarn service recovers.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18000 | Critical | Yes |

Parameters

| Name | Meaning |
|------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceNam | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

The cluster cannot provide Yarn services. Users cannot run new applications. Submitted applications cannot be run.

Possible Causes

- The ZooKeeper service is abnormal.
- The HDFS service is abnormal.
- There is no active ResourceManager instance in the Yarn cluster.
- All the NodeManagers in the Yarn cluster are abnormal.

Procedure

Check ZooKeeper service status.

Step 1 On the FusionInsight Manager, check whether the alarm list contains **ALM-13000 ZooKeeper Service Unavailable**.

- If yes, go to [Step 2](#).
- If no, go to [Step 3](#).

Step 2 Rectify the fault by following the steps provided in **ALM-13000 ZooKeeper Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

Check the HDFS service status.

Step 3 On the FusionInsight Manager, check whether the alarm list contains the HDFS alarms.

- If yes, go to [Step 4](#).
- If no, go to [Step 5](#).

Step 4 Choose **O&M > Alarm > Alarms**, handle HDFS alarms based on the alarm help, and check whether the Yarn alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check the ResourceManager status in the Yarn cluster.

Step 5 On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn**.

Step 6 In **Dashboard**, check whether there is an active ResourceManager instance in the Yarn cluster.

- If yes, go to [Step 7](#).
- If no, go to [Step 10](#).

Check the NodeManager node status in the Yarn cluster.

Step 7 On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Instance**.

Step 8 Query NodeManager **Running Status**, and check whether there are unhealthy nodes.

- If yes, go to [Step 9](#).
- If no, go to [Step 10](#).


Step 9 Rectify the fault by following the steps provided in **ALM-18002 NodeManager Heartbeat Lost** or **ALM-18003 NodeManager Unhealthy**. After the fault is rectified, check whether the Yarn alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

Collect fault information.

Step 10 On the FusionInsight Manager portal of the active cluster, choose **O&M > Log > Download**.

Step 11 Select **Yarn** in the required cluster from the **Service**.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.146 ALM-18002 NodeManager Heartbeat Lost

Description

The system checks the number of lost NodeManager nodes every 30 seconds, and compares the number with the threshold. The Number of Lost Nodes indicator has a default threshold. The alarm is generated when the value of Number of Lost Nodes exceeds the threshold.

To change the threshold, on FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Yarn**. On the displayed page, choose **Configurations > All Configurations**, and change the value of **yarn.nodemanager.lost.alarm.threshold**. You do not need to restart Yarn to make the change take effect.

The default threshold is 0. The alarm is generated when the number of lost nodes exceeds the threshold, and is cleared when the number of lost nodes is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18002 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Lost Host | Specifies the list of hosts with lost nodes. |

Impact on the System


- The lost NodeManager node cannot provide the Yarn service.
- The number of containers decreases, so the cluster performance deteriorates.

Possible Causes

- NodeManager is forcibly deleted without decommission.
- All the NodeManager instances are stopped or the NodeManager process is faulty.
- The host where the NodeManager node resides is faulty.
- The network between the NodeManager and ResourceManager is disconnected or busy.

Procedure

Check the NodeManager status.

Step 1 On the FusionInsight Manager, and choose **O&M > Alarm > Alarms**. Click  before the alarm and obtain lost nodes in **Additional Information**.

Step 2 Check whether the lost nodes are hosts that have been manually deleted without decommission.

- If yes, go to **Step 3**.

- If no, go to [Step 5](#).

Step 3 After the setting, Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn**. On the displayed page, choose **Configurations** > **All Configurations**. Search for **yarn.nodemanager.lost.alarm.threshold** and change its value to the number of hosts that are not out of service and proactively deleted. After the setting, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Step 4 Manually clear the alarm. Note that decommission must be performed before deleting hosts.

Step 5 On the FusionInsight Manager portal, choose **Cluster** > **Hosts**, and check whether the nodes obtained in [Step 1](#) are healthy.

- If yes, go to [Step 7](#).
- If no, go to [Step 6](#).

Step 6 Rectify the node fault based on **ALM-12006 Node Fault** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Check the process status.

Step 7 On the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Instance**, and check whether there are NodeManager instances whose status is not **Good**.

- If yes, go to [Step 10](#).
- If no, go to [Step 8](#).

Step 8 Check whether the NodeManager instance is deleted.

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

Step 9 Restart the active and standby ResourceManager instances, and check whether the alarm is cleared.

 **NOTE**

- Restarting the active ResourceManager instance will trigger a active/standby switchover of the ResourceManager instance. During the switchover, Yarn cannot submit new jobs, but submitted jobs are not affected. The Yarn component and components that depend on Yarn generate temporary service unavailability alarms.
- Restart the standby ResourceManager instance. Services are not affected.
- If yes, no further action is required.
- If no, go to [Step 13](#).

Check the instance status.

Step 10 Select NodeManager instances which running state is not **Normal** and restart them. Check whether the alarm is cleared.

 NOTE

During NodeManager restart, containers submitted to this node may be retried to other nodes.

- If yes, no further action is required.
- If no, go to [Step 11](#).

Check the network status.

Step 11 Log in to the management node, **ping** the IP address of the lost NodeManager node to check whether the network is disconnected or busy.

- If yes, go to [Step 12](#).
- If no, go to [Step 13](#).


Step 12 Rectify the network, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 13](#).

Collect fault information.

Step 13 On the FusionInsight Manager in the active cluster, choose **O&M > Log > Download**.

Step 14 Select **Yarn** in the required cluster from the **Service**.

Step 15 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.147 ALM-18003 NodeManager Unhealthy

Description

The system checks the number of unhealthy NodeManager nodes every 30 seconds, and compares the number with the threshold. The Unhealthy Nodes indicator has a default threshold. This alarm is generated when the value of the Unhealthy Nodes indicator exceeds the threshold.

To change the threshold, on FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Yarn**. On the displayed page, choose **Configurations > All Configurations**, and change the value of **yarn.nodemanager.unhealthy.alarm.threshold**. You do not need to restart Yarn to make the change take effect.

The default threshold is 0. The alarm is generated when the number of unhealthy nodes exceeds the threshold, and is cleared when the number of unhealthy nodes is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18003 | Major | Yes |

Parameters

| Name | Meaning |
|----------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Unhealthy Host | Specifies the list of hosts with unhealthy nodes. |

Impact on the System


- The faulty NodeManager node cannot provide the Yarn service.
- The number of containers decreases, so the cluster performance deteriorates.

Possible Causes

- The hard disk space of the host where the NodeManager node resides is insufficient.
- User **omm** does not have the permission to access a local directory on the NodeManager node.

Procedure

Check the hard disk space of the host.

Step 1 On the FusionInsight Manager, and choose **O&M > Alarm > Alarms**. Click  before the alarm and obtain unhealthy nodes in **Additional Information**.

Step 2 Choose **Cluster > Name of the desired cluster > Services > Yarn > Instance**, select the NodeManager instance corresponding to the host, choose **Instance**

Configurations > All Configurations and view disks corresponding to **yarn.nodemanager.local-dirs** and **yarn.nodemanager.log-dirs**.

Step 3 Choose **O&M > Alarm > Alarms**. In the alarm list, check whether the related disk has the alarm **ALM-12017 Insufficient Disk Capacity**.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Rectify the disk fault based on **ALM-12017 Insufficient Disk Capacity** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Step 5 Choose **Hosts > Name of the desired host**. On the **Dashboard** page, check the disk usage of the corresponding partition. Check whether the percentage of the used space of the mounted disk exceeds the value of **yarn.nodemanager.disk-health-checker.max-disk-utilization-per-disk-percentage**

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 Reduce the disk usage to less than the value of **yarn.nodemanager.disk-health-checker.max-disk-utilization-per-disk-percentage**, wait for 10 to 20 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check the access permission of the local directory on each NodeManager node.

Step 7 Obtain the NodeManager directory viewed in **Step 2**, log in to each NodeManager node as user **root**, and go to the obtained directory.

Step 8 Run the **ll** command to check whether the permission of the **localdir** and **containerlogs** folders is **755** and whether **User:Group** is **omm:ficommon**.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Run the following command to set the permission to **755** and **User:Group** to **omm:ficommon**:

```
chmod 755 <folder_name>
```

```
chown omm:ficommon <folder_name>
```


Step 10 Wait for 10 to 20 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On the FusionInsight Manager in the active cluster, choose **O&M > Log > Download**.

Step 12 Select **Yarn** in the required cluster from the **Service**.

Step 13 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.148 ALM-18008 Heap Memory Usage of ResourceManager Exceeds the Threshold

Description

The system checks the heap memory usage of Yarn ResourceManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Yarn ResourceManager exceeds the threshold (95% of the maximum memory by default).

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > Yarn** to change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the heap memory usage of Yarn ResourceManager is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the heap memory usage of Yarn ResourceManager is less than or equal to 95% of the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18008 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |

| Name | Meaning |
|-------------------|--|
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

When the heap memory usage of Yarn ResourceManager is overhigh, the performance of Yarn task submission and operation is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

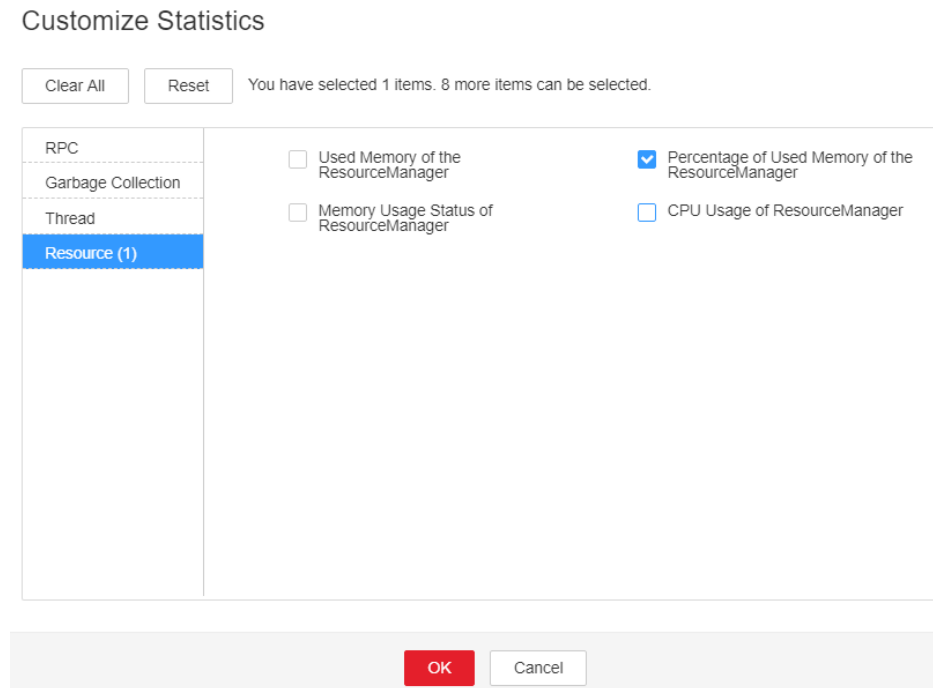
The heap memory of the Yarn ResourceManager instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the heap memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Heap Memory Usage of Yarn ResourceManager Exceeds the Threshold > Location**. Check the HostName of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Instance > ResourceManager** (Indicates the host name of the instance for which the alarm is generated). Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > ResourceManager > Percentage of Used Memory of the ResourceManager**. Check the heap memory usage.

Figure 7-92 Percentage of Used Memory of the ResourceManager



- Step 3** Check whether the used heap memory of ResourceManager reaches 95% of the maximum heap memory specified for ResourceManager.
- If yes, go to [Step 4](#).
 - If no, go to [Step 6](#).
- Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations** > **ResourceManager** > **System**. Increase the value of **GC_OPTS** parameter as required, click **Save**. Restart the role instance.

 **NOTE**

- Restarting the active ResourceManager instance will trigger a active/standby switchover of the ResourceManager instance. During the switchover, Yarn cannot submit new jobs, but submitted jobs are not affected. The Yarn component and components that depend on Yarn generate temporary service unavailability alarms.

Restart the standby ResourceManager instance. Services are not affected.

- The mapping between the number of NodeManager instances in a cluster and the memory size of ResourceManager is as follows:
 - If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
 - If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=1G
 - If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms10G -Xmx10G -XX:NewSize=1G -XX:MaxNewSize=2G
 - If the number of NodeManager instances in the cluster reaches 1000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms20G -Xmx20G -XX:NewSize=1G -XX:MaxNewSize=2G
 - If the number of NodeManager instances in the cluster reaches 2000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms40G -Xmx40G -XX:NewSize=2G -XX:MaxNewSize=4G
 - If the number of NodeManager instances in the cluster reaches 3000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms60G -Xmx60G -XX:NewSize=2G -XX:MaxNewSize=4G
 - If the number of NodeManager instances in the cluster reaches 4000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms80G -Xmx80G -XX:NewSize=2G -XX:MaxNewSize=4G
 - If the number of NodeManager instances in the cluster reaches 5000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms100G -Xmx100G -XX:NewSize=3G -XX:MaxNewSize=6G

Step 5 Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Yarn

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.149 ALM-18009 Heap Memory Usage of JobHistoryServer Exceeds the Threshold

Description

The system checks the heap memory usage of Mapreduce JobHistoryServer every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Mapreduce JobHistoryServer exceeds the threshold (95% of the maximum memory by default).

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > Mapreduce** to change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the heap memory usage of MapReduce JobHistoryServer is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the heap memory usage of MapReduce JobHistoryServer is less than or equal to 95% of the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18009 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |

| Name | Meaning |
|-------------------|--|
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

When the heap memory usage of Mapreduce JobHistoryServer is overhigh, the performance of Mapreduce log archiving is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

The heap memory of the Mapreduce JobHistoryServer instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18009 Heap Memory Usage of MapReduce JobHistoryServer Exceeds the Threshold > Location**. Check the HostName of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Mapreduce > Instance > JobHistoryServer**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > JobHistoryServer heap memory usage statistics**. JobHistoryServer indicates the corresponding HostName of the instance for which the alarm is generated. Check the heap memory usage.
- Step 3** Check whether the used heap memory of JobHistoryServer reaches 95% of the maximum heap memory specified for JobHistoryServer.
 - If yes, go to **Step 4**.
 - If no, go to **Step 6**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Mapreduce > Configurations > All Configurations > JobHistoryServer > System**. Increase the value of **GC_OPTS** parameter as required, click **Save**. Click **OK** and restart the role instance.

NOTE

- The mapping between the number of historical tasks (10000) and the memory of JobHistoryServer is as follows:
-Xms30G -Xmx30G -XX:NewSize=1G -XX:MaxNewSize=2G
- During the restart of JobHistoryServer, the status query of tasks such as Hive is affected, and the query result may be inaccurate.

Step 5 Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Mapreduce

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.150 ALM-18010 ResourceManager GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the ResourceManager process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18010 | Major | Yes |

Parameters

| Name | Meaning |
|--------|---|
| Source | Specifies the cluster for which the alarm is generated. |

| Name | Meaning |
|-------------------|--|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

A long GC duration of the ResourceManager process may interrupt the services.

Possible Causes

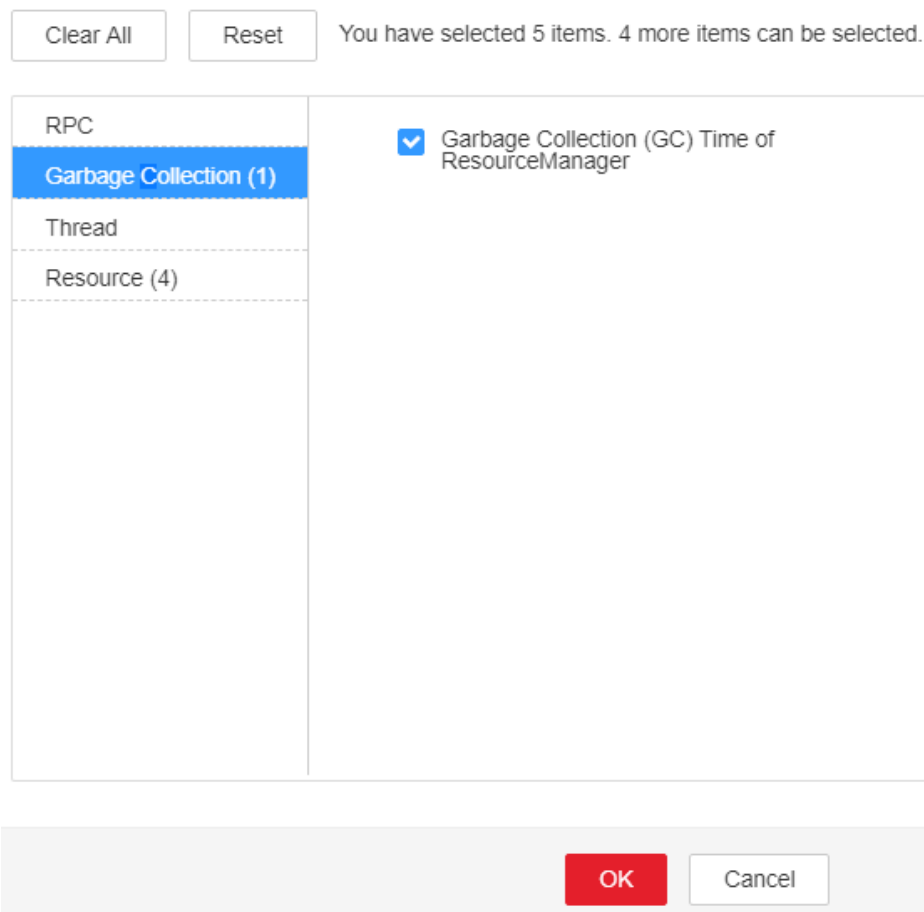
The heap memory of the ResourceManager instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18010 ResourceManager GC Time Exceeds the Threshold > Location** to check the IP address of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > *Name of the desired cluster* > Services > Yarn > Instance > ResourceManager (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Garbage Collection (GC) Time of ResourceManager** to check the GC duration statistics of the Broker process collected every minute.

Figure 7-93 Garbage Collection (GC) Time of ResourceManager
Customize Statistics



Step 3 Check whether the GC duration of the ResourceManager process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations** > **ResourceManager** > **System** to increase the value of **GC_OPTS** parameter as required.

 **NOTE**

The mapping between the number of NodeManager instances in a cluster and the memory size of ResourceManager is as follows:

- If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms10G -Xmx10G -XX:NewSize=1G -XX:MaxNewSize=2G
- If the number of NodeManager instances in the cluster reaches 1000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms20G -Xmx20G -XX:NewSize=1G -XX:MaxNewSize=2G
- If the number of NodeManager instances in the cluster reaches 2000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms40G -Xmx40G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 3000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms60G -Xmx60G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 4000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms80G -Xmx80G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 5000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms100G -Xmx100G -XX:NewSize=3G -XX:MaxNewSize=6G

Step 5 Save the configuration and restart the ResourceManager instance.

 **NOTE**

- Restarting the active ResourceManager instance will trigger a active/standby switchover of the ResourceManager instance. During the switchover, Yarn cannot submit new jobs, but submitted jobs are not affected. The Yarn component and components that depend on Yarn generate service unavailable alarms for a short period of time.
- Restart the standby ResourceManager instance. Services are not affected.


Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 8 Select **ResourceManager** in the required cluster from the **Service**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.151 ALM-18011 NodeManager GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the NodeManager process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18011 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

A long GC duration of the NodeManager process may interrupt the services.

Possible Causes

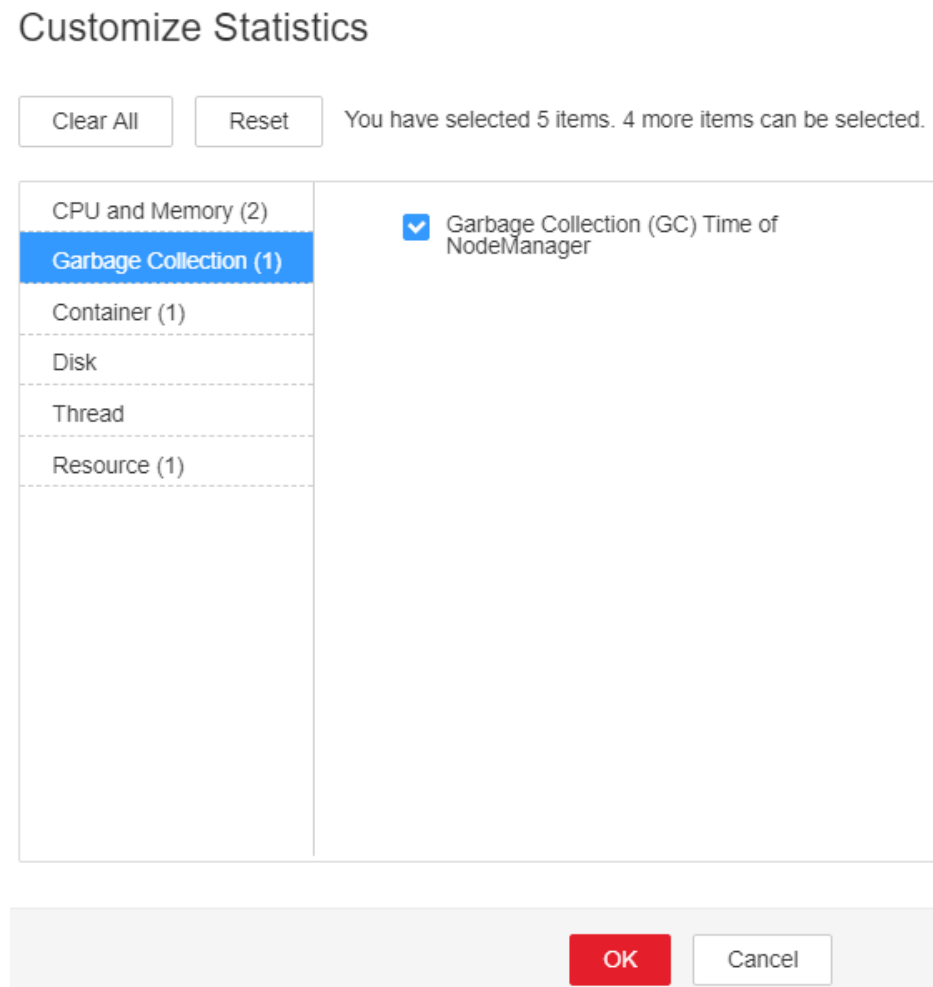
The heap memory of the NodeManager instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18011 NodeManager GC Time Exceeds the Threshold > Location** to check the IP address of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Instance > NodeManager (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Garbage Collection (GC) Time of NodeManager** to check the GC duration statistics of the Broker process collected every minute.

Figure 7-94 Garbage Collection (GC) Time of NodeManager



- Step 3** Check whether the GC duration of the NodeManager process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations** > **NodeManager** > **System** to increase the value of **GC_OPTS** parameter as required.

 **NOTE**

The mapping between the number of NodeManager instances in a cluster and the memory size of NodeManager is as follows:

- If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters for NodeManager instances are as follows: -Xms2G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters for NodeManager instances are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters for NodeManager instances are as follows: -Xms8G -Xmx8G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Save the configuration and restart the NodeManager instance.

 **NOTE**

During NodeManager restart, containers submitted to this node may be retried to other nodes.


Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **NodeManager** in the required cluster from the **Service**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.152 ALM-18012 JobHistoryServer GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the JobHistoryServer process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18012 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

A long GC duration of the JobHistoryServer process may interrupt the services.

Possible Causes

The heap memory of the JobHistoryServer instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18012 JobHistoryServer GC Time Exceeds the Threshold > Location** to check the IP address of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > MapReduce > Instance > JobHistoryServer (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Garbage Collection (GC) Time of the JobHistoryServer** to check the GC duration statistics of the Broker process collected every minute.
- Step 3** Check whether the GC duration of the JobHistoryServer process collected every minute exceeds the threshold (12 seconds by default).
- If yes, go to **Step 4**.
 - If no, go to **Step 7**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > MapReduce > Configurations > All Configurations > JobHistoryServer > System** to increase the value of **GC_OPTS** parameter as required.

 **NOTE**

The mapping between the number of historical tasks (10000) and the memory of the JobHistoryServer is as follows:

```
-Xms30G -Xmx30G -XX:NewSize=1G -XX:MaxNewSize=2G
```

- Step 5** Save the configuration and restart the JobHistoryServer instance.

 **NOTE**


During the restart of JobHistoryServer, the status query of tasks such as Hive is affected, and the query result may be inaccurate.

- Step 6** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to **Step 7**.

Collect fault information.

- Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

- Step 8** Select **JobHistoryServer** in the required cluster from the **Service**.

- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

- Step 10** Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.153 ALM-18013 ResourceManager Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of ResourceManager every 30 seconds. This alarm is generated when the direct memory usage of ResourceManager instances exceeds the threshold (90% of the maximum memory).

This alarm is automatically cleared when the direct memory usage is less than the threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 18013 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

If the available direct memory of ResourceManager is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

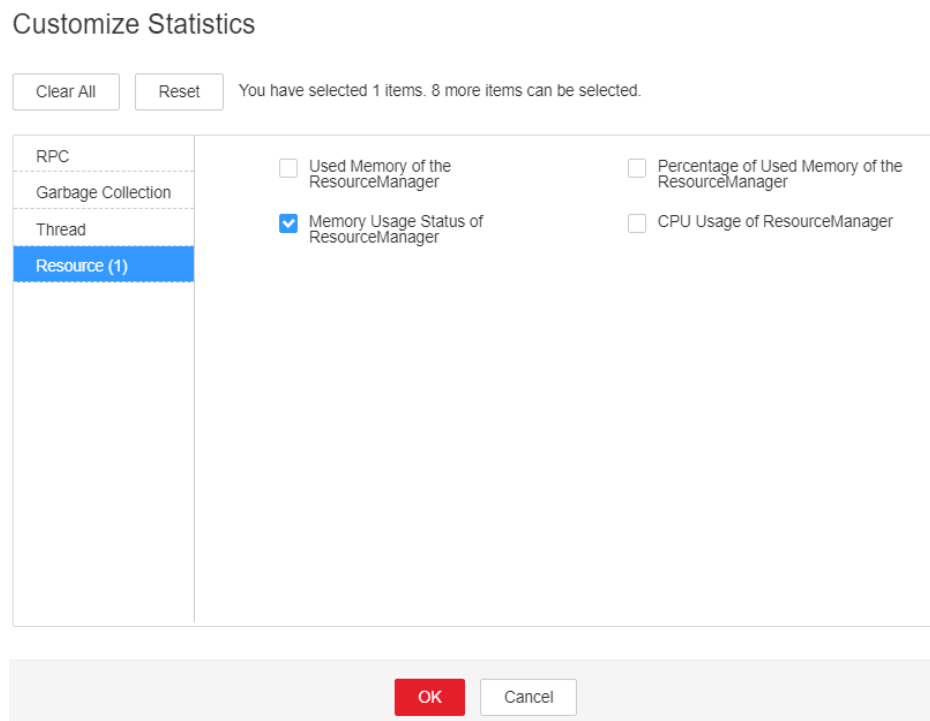
The direct memory of ResourceManager instances is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-18013 ResourceManager Direct Memory Usage Exceeds the Threshold > Location**. View the IP address of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Yarn**. On the page that is displayed, click the **Instances** tab and click the ResourceManager instance for which this alarm is generated. Click the drop-down list in the upper right corner of the chart area, choose **Customize > Resource**, and select **Memory Usage Status of ResourceManager** to check the direct memory usage.

Figure 7-95 Customizing ResourceManager memory usage details



- Step 3** Check whether the used direct memory of a ResourceManager instance reaches 90% (default threshold) of the maximum direct memory allocated to it.
- If yes, go to **Step 4**.
 - If no, go to **Step 9**.
- Step 4** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Yarn > Configurations > All Configurations > ResourceManager > System**. Check whether **-XX:MaxDirectMemorySize** exists in the **GC_OPTS** parameter.
- If yes, go to **Step 5**.
 - If no, go to **Step 7**.
- Step 5** Delete the **-XX:MaxDirectMemorySize** parameter from **GC_OPTS** and save the configuration.

 NOTE

MaxDirectMemorySize indicates the maximum off-heap memory size. If the **MaxDirectMemorySize** parameter of ResourceManager is not specified, the memory of ResourceManager is not limited. By default, **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter is not set.

Step 6 Perform the following steps to restart the ResourceManager instance:

NOTICE

- Restarting the standby ResourceManager instance does not affect services.
- During the ResourceManager switchover, new jobs cannot be submitted to Yarn, but submitted jobs are not affected.

-
1. On the Yarn service page, click the **Instances** tab, select the **ResourceManager (Standby)** instance, choose **More**, select **Restart Instance**, and verify the password to restart the instance.
 2. After the standby instance is restarted, click the **Dashboard** tab of Yarn, choose **More**, select **Perform ResourceManager Switchover**, and verify the password to perform an active/standby switchover.
 3. After the active/standby switchover is complete, click the **Instances** tab on the Yarn service page, select the **ResourceManager (Standby)** instance, choose **More**, select **Restart Instance**, and verify the password to restart the instance. Wait until the instance is restarted.

Step 7 Check whether **ALM-18008 Heap Memory Usage of ResourceManager Exceeds the Threshold** exists.

- If yes, rectify the fault by referring to **ALM-18008 Heap Memory Usage of ResourceManager Exceeds the Threshold**.
- If no, go to [Step 8](#).


Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 10 Expand the **Service** drop-down list, and select **ResourceManager** for the target cluster.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.154 ALM-18014 NodeManager Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the Yarn service every 30 seconds. This alarm is generated when the direct memory usage of a NodeManager instance exceeds the threshold (90% of the maximum memory).

The alarm is cleared when the direct memory usage is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18014 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

If the available direct memory of the Yarn service is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

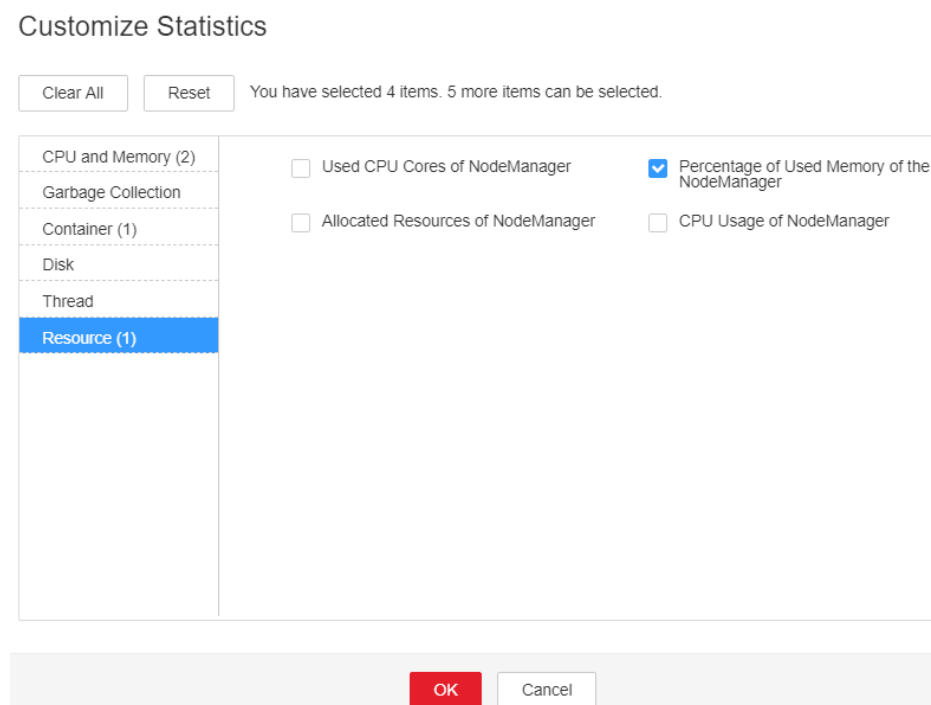
The direct memory of the NodeManager instance is overused or the direct memory is inappropriately allocated.

Procedure

Check the direct memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18014 NodeManager Direct Memory Usage Exceeds the Threshold > Location** to check the IP address of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Instance > NodeManager (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Resource > Percentage of Used Memory of the NodeManager** to check the direct memory usage.

Figure 7-96 Percentage of Used Memory of the NodeManager



- Step 3** Check whether the used direct memory of NodeManager reaches 90% of the maximum direct memory specified for NodeManager by default.
- If yes, go to **Step 4**.
 - If no, go to **Step 9**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Configurations > All Configurations > NodeManager > System** to check whether "-XX:MaxDirectMemorySize" exists in the **GC_OPTS** parameter.
- If yes, go to **Step 5**.

- If no, go to [Step 7](#).

Step 5 In the **GC_OPTS** parameter, delete "-XX:MaxDirectMemorySize".

Step 6 Save the configuration and restart the NodeManager instance.

 **NOTE**

During NodeManager restart, containers submitted to this node may be retried to other nodes.

Step 7 Check whether the **ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold** exists.

- If yes, handle the alarm by referring to **ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold**.
- If no, go to [Step 8](#).


Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 10 Select **NodeManager** in the required cluster from the **Service**.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.155 ALM-18015 JobHistoryServer Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the MapReduce service every 30 seconds. This alarm is generated when the direct memory usage of a JobHistoryServer instance exceeds the threshold (90% of the maximum memory).

The alarm is cleared when the direct memory usage is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18015 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

If the available direct memory of the MapReduce service is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the JobHistoryServer instance is overused or the direct memory is inappropriately allocated.

Procedure


Check the direct memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18015 JobHistoryServer Direct Memory Usage Exceeds the Threshold > Location** to check the IP address of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > MapReduce > Instance > JobHistoryServer (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Memory Usage Status of JobHistoryServer** to check the direct memory usage.

- Step 3** Check whether the used direct memory of JobHistoryServer reaches 90% of the maximum direct memory specified for JobHistoryServer by default.
- If yes, go to [Step 4](#).
 - If no, go to [Step 9](#).
- Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **MapReduce** > **Configurations** > **All Configurations** > **JobHistoryServer** > **System** to check whether "-XX:MaxDirectMemorySize" exists in the **GC_OPTS** parameter.
- If yes, go to [Step 5](#).
 - If no, go to [Step 7](#).
- Step 5** In the **GC_OPTS** parameter, delete "-XX:MaxDirectMemorySize".
- Step 6** Save the configuration and restart the JobHistoryServer instance.

 **NOTE**

During the restart of JobHistoryServer, the status query of tasks such as Hive is affected, and the query result may be inaccurate.

- Step 7** Check whether the **ALM-18009 Heap Memory Usage of JobHistoryServer Exceeds the Threshold** exists.
- If yes, handle the alarm by referring to **ALM-18009 Heap Memory Usage of JobHistoryServer Exceeds the Threshold**.
 - If no, go to [Step 8](#).
- Step 8** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 9](#).
- Collect fault information.**
- Step 9** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.
- Step 10** Select **JobHistoryServer** in the required cluster from the **Service**.
- Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 12** Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.156 ALM-18016 Non Heap Memory Usage of ResourceManager Exceeds the Threshold

Description

The system checks the Non Heap memory usage of Yarn ResourceManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the Non Heap memory usage of Yarn ResourceManager exceeds the threshold (90% of the maximum memory by default).

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > Yarn** to change the threshold.

The alarm is cleared when the Non Heap memory usage is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18016 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

When the Non Heap memory usage of Yarn ResourceManager is overhigh, the performance of Yarn task submission and operation is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

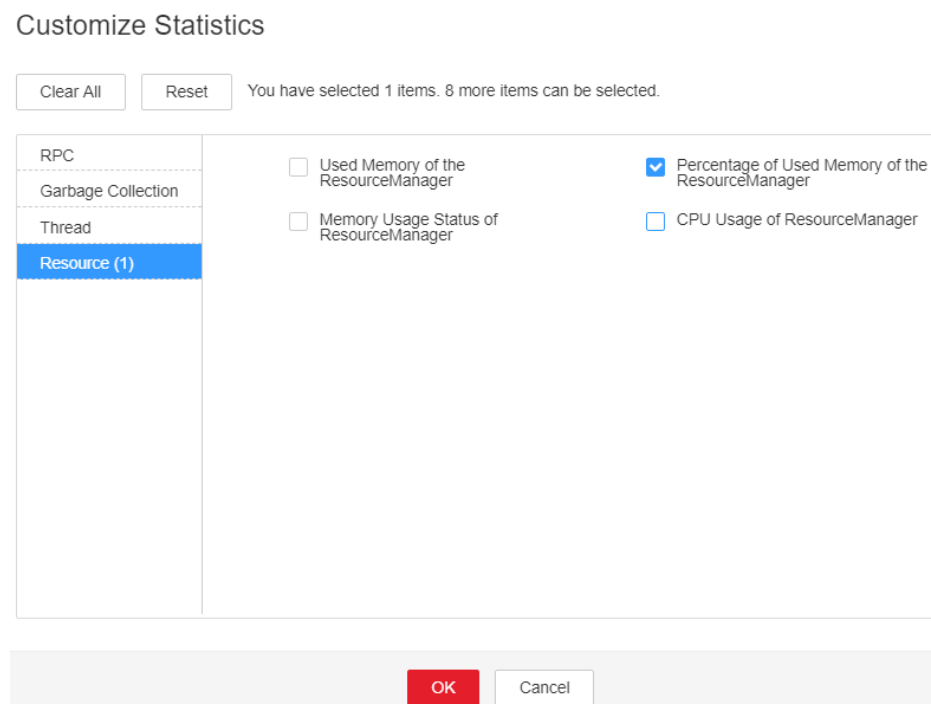
The Non Heap memory of the Yarn ResourceManager instance on the node is overused or the Non Heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the Non Heap memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18016 Non Heap Memory Usage of Yarn ResourceManager Exceeds the Threshold > Location**. Check the HostName of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Instance > ResourceManager**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Percentage of Used Memory of the ResourceManager**. ResourceManager indicates the corresponding HostName of the instance for which the alarm is generated. Check the Non Heap memory usage.

Figure 7-97 Percentage of Used Memory of the ResourceManage



- Step 3** Check whether the used Non Heap memory of ResourceManager reaches 90% of the maximum Non Heap memory specified for ResourceManager by default.
- If yes, go to **Step 4**.
 - If no, go to **Step 6**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Configurations > All Configurations > ResourceManager > System**. Adjust the **GC_OPTS** memory parameter of

ResourceManager. Save the configuration and restart the ResourceManager instance.

 **NOTE**

- Restarting the active ResourceManager instance will trigger a active/standby switchover of the ResourceManager instance. During the switchover, Yarn cannot submit new jobs, but submitted jobs are not affected. The Yarn component and components that depend on Yarn generate service unavailable alarms for a short period of time.

Restart the standby ResourceManager instance. Services are not affected.

- The mapping between the number of NodeManager instances in a cluster and the memory size of ResourceManager is as follows:
 - If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
 - If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=1G
 - If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms10G -Xmx10G -XX:NewSize=1G -XX:MaxNewSize=2G
 - If the number of NodeManager instances in the cluster reaches 1000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms20G -Xmx20G -XX:NewSize=1G -XX:MaxNewSize=2G
 - If the number of NodeManager instances in the cluster reaches 2000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms40G -Xmx40G -XX:NewSize=2G -XX:MaxNewSize=4G
 - If the number of NodeManager instances in the cluster reaches 3000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms60G -Xmx60G -XX:NewSize=2G -XX:MaxNewSize=4G
 - If the number of NodeManager instances in the cluster reaches 4000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms80G -Xmx80G -XX:NewSize=2G -XX:MaxNewSize=4G
 - If the number of NodeManager instances in the cluster reaches 5000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms100G -Xmx100G -XX:NewSize=3G -XX:MaxNewSize=6G

Step 5 Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Yarn

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.157 ALM-18017 Non Heap Memory Usage of NodeManager Exceeds the Threshold

Description

The system checks the Non Heap memory usage of Yarn NodeManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the Non Heap memory usage of Yarn NodeManager exceeds the threshold (90% of the maximum memory by default).

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > Yarn** to change the threshold.

The alarm is cleared when the Non Heap memory usage is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18017 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

When the Non Heap memory usage of Yarn NodeManager is overhigh, the performance of Yarn task submission and operation is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes


The Non Heap memory of the Yarn NodeManager instance on the node is overused or the Non Heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the Non Heap memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18017 Non Heap Memory Usage of Yarn NodeManager Exceeds the Threshold > Location**. Check the HostName of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Instance > NodeManager**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Resource > Percentage of Used Memory of the NodeManager**. NodeManager indicates the corresponding HostName of the instance for which the alarm is generated. Check the Non Heap memory usage.

Figure 7-98 Percentage of Used Memory of the NodeManager



- Step 3** Check whether the used Non Heap memory of NodeManager reaches 90% of the maximum Non Heap memory specified for NodeManager by default.

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations** > **NodeManager** > **System**. Adjust the **GC_OPTS** memory parameter of NodeManager, click **Save**, and click **OK**, and restart the role instance.

 **NOTE**

- During NodeManager restart, containers submitted to this node may be retried to other nodes.
- The mapping between the number of NodeManager instances in a cluster and the memory size of NodeManager is as follows:
 - If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters for NodeManager instances are as follows: -Xms2G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
 - If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters for NodeManager instances are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
 - If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters for NodeManager instances are as follows: -Xms8G -Xmx8G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Yarn

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.158 ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold

Description

The system checks the heap memory usage of Yarn NodeManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Yarn NodeManager exceeds the threshold (95% of the maximum memory by default).

The alarm is cleared when the heap memory usage is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18018 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

When the heap memory usage of Yarn NodeManager is overhigh, the performance of Yarn task submission and operation is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes


The heap memory of the Yarn NodeManager instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the heap memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold > Location**. Check the HostName of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Instance > NodeManager**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > Resource > Percentage of Used Memory of the NodeManager** to check the heap memory usage.

Figure 7-99 Percentage of Used Memory of the NodeManager



- Step 3** Check whether the used heap memory of NodeManager reaches 95% of the maximum heap memory specified for NodeManager.
- If yes, go to **Step 4**.
 - If no, go to **Step 6**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Yarn > Configurations > All Configurations > NodeManager > System**. Increase the value of **GC_OPTS** parameter as required, click **Save**, and click **OK**, and restart the role instance.

 NOTE

- During NodeManager restart, containers submitted to this node may be retried to other nodes.
- The mapping between the number of NodeManager instances in a cluster and the memory size of NodeManager is as follows:
 - If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters for NodeManager instances are as follows: -Xms2G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
 - If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters for NodeManager instances are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
 - If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters for NodeManager instances are as follows: -Xms8G -Xmx8G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Yarn

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.159 ALM-18019 Non Heap Memory Usage of JobHistoryServer Exceeds the Threshold

Description

The system checks the Non Heap memory usage of MapReduce JobHistoryServer every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the Non Heap memory usage of MapReduce JobHistoryServer exceeds the threshold (90% of the maximum memory by default).

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > MapReduce** to change the threshold.

The alarm is cleared when the Non Heap memory usage is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18019 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

When the Non Heap memory usage of MapReduce JobHistoryServer is overhigh, the performance of MapReduce task submission and operation is affected. In addition, a memory overflow may occur so that the MapReduce service is unavailable.

Possible Causes

The Non Heap memory of the MapReduce JobHistoryServer instance on the node is overused or the Non Heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the Non Heap memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > ALM-18019 Non Heap Memory Usage of MapReduce JobHistoryServer Exceeds**

the Threshold > Location. Check the HostName of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > MapReduce > Instance > JobHistoryServer**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize > JobHistoryServer Non Heap memory usage statistics**. JobHistoryServer indicates the corresponding HostName of the instance for which the alarm is generated. Check the Non Heap memory usage.

Step 3 Check whether the used Non Heap memory of JobHistoryServer reaches 90% of the maximum Non Heap memory specified for JobHistoryServer.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > MapReduce > Configurations > All Configurations > JobHistoryServer > System**. Adjust the **GC_OPTS** memory parameter of the NodeManager, click **Save**, and click **OK**, and restart the role instance.

NOTE

- During the restart of JobHistoryServer, the status query of tasks such as Hive is affected, and the query result may be inaccurate.
- The mapping between the number of historical tasks (10000) and the memory of the JobHistoryServer is as follows:
-Xms30G -Xmx30G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- MapReduce

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.160 ALM-18020 Yarn Task Execution Timeout

Alarm Description

The system checks MapReduce and Spark tasks (except for permanent JDBC tasks) submitted to Yarn every 15 minutes. This alarm is generated when the task execution time exceeds the timeout duration specified by the user. However, the task can be properly executed. The client timeout parameter of MapReduce is `mapreduce.application.timeout.alarm` and that of Spark is `spark.application.timeout.alarm`. The unit is ms.

This alarm is cleared when the task is finished or terminated.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 18020 | Minor | Yes |

Alarm Parameters

| Parameter | Description |
|-------------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| ApplicationName | Specifies the object (application ID) for which the alarm was generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

The alarm persists after task execution times out. However, the task can still be properly executed, so this alarm does not exert any impact on the system.

Possible Causes

- The specified timeout duration is shorter than the required execution time.
- The queue resources for task running are insufficient.
- Task data skew occurs. As a result, some tasks process a large amount of data and take a long time to execute.


Handling Procedure

Check whether the timeout interval is correctly set.

- Step 1** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. The **Alarms** page is displayed.
- Step 2** Select the alarm whose ID is **18020**. In the alarm details, view **Location** to obtain the timeout task name and timeout duration.
- Step 3** Based on the task name and timeout interval, choose **Cluster > Name of the desired cluster > Services > Yarn > ResourceManager (Active)** to log in to the native Yarn page. Then find the task on the native page, check its **StartTime** and calculate the task execution time based on the current system time. Check whether the task execution time exceeds the timeout duration.
- If yes, go to **Step 4**.
 - If no, go to **Step 10**.
- Step 4** Evaluate the expected task execution time based on the service and compare it with the task timeout interval. If the timeout interval is too short, set the timeout interval (**mapreduce.application.timeout.alarm** or **spark.application.timeout.alarm**) of the client to the task expected execution time. Run the task again and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to **Step 5**.

Check whether the queue resources are sufficient.

- Step 5** Find the task on the native page and view the queue name of the task. Click **Scheduler** on the left of the native page. On the **Applications Queues** page, find the corresponding queue name and expand the queue details, as shown in the following figure.



- Step 6** Check whether the value of **Used Resources** in the queue details is approximately equal to the value of **Max Resources**, which indicates that the resources in the queue submitted by the task have been used up. If the queue resources are insufficient, choose **Tenant Resources > Dynamic Resource Plan > Resource Distribution Policy** on FusionInsight Manager and increase the value of **Max Resources** for the queue. Run the task again and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to **Step 7**.

Check whether data skew occurs.

Step 7 On the native Yarn page, click *task ID* (for example, **application_1565337919723_0002**) > **Tracking URL:ApplicationMaster** > **job_1565337919723_0002**. The following page is displayed.

| Attempt Number | Start Time | Node | Logs |
|----------------|-------------------------------|--------------------|------|
| 1 | Fri Aug 9 17:23:05 +0800 2019 | 187-7-66-181:26010 | logs |

| Task Type | Progress | Total | Pending | Running | Complete |
|-----------|----------|-------|---------|---------|----------|
| Map | 10 | 0 | 10 | 0 | 0 |
| Reduce | 1 | 1 | 0 | 0 | 0 |

| Attempt Type | New | Running | Failed | Killed | Successful |
|--------------|-----|---------|--------|--------|------------|
| Maps | 0 | 10 | 0 | 0 | 0 |
| Reduces | 1 | 0 | 0 | 0 | 0 |

Step 8 Choose **Job > Map tasks** or **Job > Reduce tasks** on the left and check whether the execution time of each Map or Reduce task differs greatly. If yes, task data skew occurs. In this case, you need to balance the task data.


Step 9 Rectify the fault based on the preceding causes and perform the tasks again. Then, check whether the alarm persists.

- If yes, go to **Step 10**.
- If no, no further action is required.

Collect the fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 11 Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.161 ALM-18021 Mapreduce Service Unavailable

Description

The alarm module checks the MapReduce service status every 60 seconds. This alarm is generated when the system detects that the MapReduce service is unavailable.

The alarm is cleared when the MapReduce service recovers.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18021 | Critical | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

The cluster cannot provide the MapReduce service. For example, MapReduce cannot be used to view task logs or the log archive function is unavailable.

Possible Causes

- The JobHistoryServer instance is abnormal.
- The KrbServer service is abnormal.
- The ZooKeeper service abnormal.
- The HDFS service abnormal.
- The Yarn service is abnormal.

Procedure

Check MapReduce service JobHistoryServer instance status.

Step 1 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **MapReduce** > **Instance**.

Step 2 Check whether the running status of JobHistoryServer is **Normal**.

- If yes, go to [Step 11](#).
- If no, go to [Step 3](#).

Check the KrbServer service status.

Step 3 In the alarm list on FusionInsight Manager, check whether **ALM-25500 KrbServer Service Unavailable** exists.

- If yes, go to [Step 4](#).
- If no, go to [Step 5](#).

Step 4 Rectify the fault by following the steps provided in **ALM-25500 KrbServer Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check the ZooKeeper service.

Step 5 In the alarm list on FusionInsight Manager, check whether **ALM-13000 ZooKeeper Service Unavailable** exists.

- If yes, go to [Step 6](#).
- If no, go to [Step 7](#).

Step 6 Rectify the fault by following the steps provided in **ALM-13000 ZooKeeper Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Check the HDFS service status.

Step 7 In the alarm list on FusionInsight Manager, check whether **ALM-14000 HDFS Service Unavailable** exists.

- If yes, go to [Step 8](#).
- If no, go to [Step 9](#).

Step 8 Rectify the fault by following the steps provided in **ALM-14000 HDFS Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Check the Yarn service status.

Step 9 In the alarm list on FusionInsight Manager, check whether **ALM-18000 Yarn Service Unavailable** exists.

- If yes, go to [Step 10](#)
- If no, go to [Step 11](#).


Step 10 Rectify the fault by following the steps provided in **ALM-18000 Yarn Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

Collect fault information.

Step 11 On the FusionInsight Manager home page of the active cluster, choose **O&M > Log > Download**.

Step 12 Select **MapReduce** in the required cluster from the **Service**.

Step 13 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.162 ALM-18022 Insufficient Yarn Queue Resources

Description

The alarm module checks Yarn queue resources every 60 seconds. This alarm is generated when available resources or ApplicationMaster (AM) resources of a queue are insufficient.

This alarm is cleared when available resources are sufficient.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18022 | Minor | Yes |

Parameters

| Parameter Name | Description |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| QueueName | Specifies the queue for which the alarm is generated. |
| QueueMetric | Specifies the metric of the queue for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

- An application being executed takes longer time.
- An application fails to be executed for a long time after being submitted.

Possible Causes

- NodeManager node resources are insufficient.
- The configured maximum resource capacity of the queue is excessively small.
- The configured maximum AM resource percentage is excessively small.

Procedure

View alarm details.

Step 1 On the FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 2 View location information of this alarm and check whether **QueueName** is **root** and **QueueMetric** is **Memory** or **QueueName** is **root** and **QueueMetric** is **vCores**.

- If yes, go to [Step 3](#).
- If no, go to [Step 4](#).

Step 3 The memory or CPU of the Yarn cluster is insufficient. In this case, log in to the node where NodeManager resides and run the **free -g** and **cat /proc/cpuinfo** commands to query the available memory and available CPU of the node, respectively. On FusionInsight Manager, increase the values of **yarn.nodemanager.resource.memory-mb** and **yarn.nodemanager.resource.cpu-vcores** for the Yarn NodeManager based on the query results. Then, restart the NodeManager instance. Check whether the alarm is cleared.

NOTE

During NodeManager restart, containers submitted to this node may be retried to other nodes.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Step 4 View location information of this alarm and check whether **QueueName** is **<Tenant Queue>** and **QueueMetric** is **Memory**, or **QueueName** is **<Tenant Queue>** and **QueueMetric** is **vCores** in **Location**, check whether **available Memory** = or **available vCores** = are included in **Additional Information**.

- If yes, go to [Step 5](#).
- If no, go to [Step 7](#).

Step 5 The memory or CPU of the tenant queue is insufficient. In this case, choose **Tenant Resources > Dynamic Resource Plan > Resource Distribution Policy** and increase the value of **Maximum Capacity**. Then, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Step 6 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations**. Enter the keyword "threshold" and click **ResourceManager**. Adjust the threshold values of the following parameters:

If **Additional Information** contains **available Memory** =, change the value of **yarn.queue.memory.alarm.threshold** to a value smaller than that of **available Memory** = in **Additional Information**.

If **Additional Information** contains **available vCores** =, change the value of **yarn.queue.vcore.alarm.threshold** to a value smaller than that of **available vCores** = in **Additional Information**.

Wait for five minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Step 7 If **available AmMemory** = or **available AmvCores** = is included in **Additional Information**, ApplicationMaster memory or CPU of the tenant queue is insufficient. In this case, choose **Tenant Resources** > **Dynamic Resource Plan** > **Queue Configuration** and increase the value of **Maximum Am Resource Percent**. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

Step 8 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations**. Enter the keyword "threshold" and click **ResourceManager**. Adjust the threshold values of the following parameters:

If **Additional Information** contains **available AmMemory** =, change the value of **yarn.queue.memory.alarm.threshold** to a value smaller than that of **available AmMemory** = in **Additional Information**.

If **Additional Information** contains **available AmvCores** =, change the value of **yarn.queue.vcore.alarm.threshold** to a value smaller than that of **available AmvCores** = in **Additional Information**.


Wait for five minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Collect fault information.

Step 9 Log in to FusionInsight Manager of the active cluster, and choose **O&M** > **Log** > **Download**.

Step 10 Select **Yarn** in the required cluster from the **Service**.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Reference

None

7.12.163 ALM-18023 Number of Pending Yarn Tasks Exceeds the Threshold

Alarm Description

The alarm module checks the number of pending applications in the Yarn root queue every 60 seconds. The alarm is generated when the number exceeds 60.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 18023 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm was generated. |
| QueueName | Specifies the queue for which the alarm was generated. |
| QueueMetric | Specifies the queue metric for which the alarm was generated. |

Impact on the System

- It takes long time to end an application.
- A new application cannot run after submission.

Possible Causes

- NodeManager node resources are insufficient.
- The maximum resource capacity of the queue and the maximum AM resource percentage are too small.
- The monitoring threshold is too small.

Handling Procedure

Check NodeManager resources.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **ResourceManager (Active)** to access the ResourceManager web UI.

Step 2 Click **Scheduler** and check whether the root queue resources are used up in **Application Queues**.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Expand the capacity of the NodeManager instance of the Yarn service. After the capacity expansion, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check the maximum queue resource capacity and the maximum AM resource percentage.

Step 4 Check whether the resources of the queue corresponding to the pending task are used up.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, choose **Tenant Resources** > **Dynamic Resource Plan** and add resources as required. Check whether the alarms are cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Adjust the monitoring thresholds.

Step 6 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Yarn** > **Applications** > **Pending Applications**, and increase the thresholds as required.


Step 7 Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 10 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.164 ALM-18024 Pending Yarn Memory Usage Exceeds the Threshold

Alarm Description

The alarm module checks the pending memory of Yarn every 60 seconds. The alarm is generated when the pending memory exceeds the threshold. Pending memory indicates the total memory that is not allocated to submitted Yarn applications.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 18024 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm was generated. |
| QueueName | Specifies the queue for which the alarm was generated. |
| QueueMetric | Specifies the queue metric for which the alarm was generated. |

Impact on the System

- It takes long time to end an application.
- A new application cannot run after submission.

Possible Causes

- NodeManager node resources are insufficient.
- The maximum resource capacity of the queue and the maximum AM resource percentage are too small.
- The monitoring threshold is too small.

Handling Procedure

Check NodeManager resources.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **ResourceManager (Active)** to access the ResourceManager web UI.

Step 2 Click **Scheduler** and check whether the root queue resources are used up in **Application Queues**.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Expand the capacity of the NodeManager instance of the Yarn service. After the capacity expansion, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check the maximum queue resource capacity and the maximum AM resource percentage.

Step 4 Check whether the resources of the queue corresponding to the pending task are used up.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, choose **Tenant Resources** > **Dynamic Resource Plan** and add resources as required. Check whether the alarms are cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Adjust the monitoring thresholds.

Step 6 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Yarn** > **CPU and Memory** > **Pending Memory**, and increase the threshold as required.


Step 7 Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 10 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.165 ALM-18025 Number of Terminated Yarn Tasks Exceeds the Threshold

Description

The alarm module checks the number of terminated applications in the Yarn root queue every 60 seconds. The alarm is generated when the number exceeds 50 for three consecutive times.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18025 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Cluster Name | Specifies the cluster for which the alarm is generated. |
| Service Name | Specifies the service for which the alarm is generated. |
| Role Name | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

A large number of application tasks are forcibly terminated.

Possible Causes

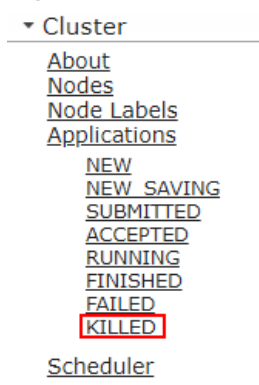
- The user forcibly terminates a large number of tasks.
- The system terminates tasks due to some error.

Procedure

Check the alarm details.


- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** to go to the alarm page.
- Step 2** View **Additional Information** in the alarm details to check whether the alarm threshold is too small.
- If yes, go to **Step 3**.
 - If no, go to **Step 4**.
- Step 3** Choose **O&M > Alarm > Thresholds > Name of the desired cluster > Yarn > Other > Terminated Applications of root queue** to modify the threshold. Go to **Step 6**.
- Step 4** Choose **Cluster > Name of the desired cluster > Services > Yarn > ResourceManager(Active)** to access the ResourceManager web UI.
- Step 5** Click **KILLED** in **Applications** and click the task on the top. View the description of **Diagnostics** and rectify the fault based on the task termination details (for example, the task is terminated by a user).

Figure 7-100 Click **KILLED**



- Step 6** Wait for 3 minutes and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to **Step 7**.

Collect the fault information.

- Step 7** On the FusionInsight Manager, choose **O&M > Log > Download**.
- Step 8** Expand the **Service** drop-down list, and select **Yarn** for the target cluster.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.166 ALM-18026 Number of Failed Yarn Tasks Exceeds the Threshold

Description

The alarm module checks the number of failed applications in the Yarn root queue every 60 seconds. The alarm is generated when the number exceeds 50 for three consecutive times.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18026 | Major | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Cluster Name | Specifies the cluster for which the alarm is generated. |
| Service Name | Specifies the service for which the alarm is generated. |
| Role Name | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

- A large number of application tasks fail to be executed.
- Failed tasks need to be submitted again.

Possible Causes


The task fails to be executed due to some error.

Procedure

Check the alarm details.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** to go to the alarm page.
- Step 2** View **Additional Information** in the alarm details to check whether the alarm threshold is too small.
 - If yes, go to **Step 3**.
 - If no, go to **Step 4**.
- Step 3** Choose **O&M > Alarm > Thresholds > Name of the desired cluster > Yarn > Other > Failed Applications of root queue** to modify the threshold. Go to **Step 6**.
- Step 4** Choose **Cluster > Name of the desired cluster > Services > Yarn > ResourceManager(Active)** to access the ResourceManager web UI.
- Step 5** Click **FAILED** in **Applications** and click the task on the top. View the description of **Diagnostics** and rectify the fault based on the task failure causes.


Figure 7-101 Click FAILED



- Step 6** Wait for 3 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 7**.

Collect the fault information.

- Step 7** On the FusionInsight Manager, choose **O&M > Log > Download**.
- Step 8** Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.167 ALM-19000 HBase Service Unavailable

Alarm Description

The alarm module checks the HBase service status every 120 seconds. This alarm is generated when the HBase service is unavailable.

This alarm is cleared when the HBase service recovers.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19000 | Critical | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |

Impact on the System

Operations such as data read/write and table creation cannot be performed.

Possible Causes

- ZooKeeper is abnormal.
- HDFS is abnormal.
- HBase is abnormal.
- The network connection is abnormal.
- The service configuration value is incorrect.

Handling Procedure

Check the ZooKeeper service status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of ZooKeeper is **Normal**.

- If yes, go to [Step 5](#).
- If no, go to [Step 2](#).

Step 2 In the alarm list, check whether **ALM-13000 ZooKeeper Service Unavailable** exists.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Rectify the fault by performing the operations provided for **ALM-13000 ZooKeeper Service Unavailable**.

Step 4 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check the HDFS service status.

Step 5 In the alarm list, check whether **ALM-14000 HDFS Service Unavailable** exists.

- If yes, go to [Step 6](#).
- If no, go to [Step 8](#).

Step 6 Rectify the fault by performing the operations provided for **ALM-14000 HDFS Service Unavailable**.

Step 7 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

Step 8 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services > HDFS**, and check whether **Safe Mode** of HDFS is **ON**.

- If yes, go to [Step 9](#).
- If no, go to [Step 12](#).

Step 9 Log in to the HDFS client as user **root**. Run the **cd** command to go to the client installation directory and run the **source bigdata_env** command.

If the cluster uses the security mode, perform security authentication. Obtain the password of user **hdfs** from the MRS cluster administrator, run the **kinit hdfs** command, and enter the password as prompted.

Step 10 Run the following command to manually exit the safe mode:

```
hdfs dfsadmin -safemode leave
```

Step 11 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

Check the HBase service status.

Step 12 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > HBase**.

Step 13 Check whether there is one active HMaster and one standby HMaster.

- If yes, go to [Step 15](#).
- If no, go to [Step 14](#).

Step 14 Click **Instances** and select the HMaster instance whose status is not **Active**. Click **More** and select **Restart Instance** to restart HMaster. Then check whether there is one active HMaster and one standby HMaster.

- If yes, go to [Step 15](#).
- If no, go to [Step 21](#).

NOTICE

During the HMaster restart, table operations cannot be performed, and the HBase web UI is inaccessible. Data read and write operations are not affected.

Step 15 Choose **Cluster**, click the name of the desired cluster, choose **Services > HBase**, and click **HMaster(Active)** to access the HMaster web UI.

 **NOTE**

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 16 Check whether at least one RegionServer exists under **Region Servers**.

- If yes, go to [Step 17](#).
- If no, go to [Step 21](#).

Step 17 Choose **Tables > System Tables** and check whether **hbase:meta**, **hbase:namespace**, and **hbase:acl** exist in the **Table Name** column, as shown in [Figure 7-102](#).

- If yes, go to [Step 18](#).
- If no, go to [Step 19](#).

Figure 7-102 HBase system tables

| Table Name | Description |
|-----------------|--|
| hbase:acl | The hbase:acl table holds information about acl. |
| hbase:hindex | The hbase:hindex table holds information about table indices. |
| hbase:meta | The hbase:meta table holds references to all User Table regions. |
| hbase:namespace | The hbase:namespace table holds information about namespaces. |

Step 18 Click **hbase:meta**, **hbase:namespace**, and **hbase:acl** to check whether all pages can be opened. If all of them can be opened, the tables are normal.

- If yes, go to [Step 19](#).
- If no, go to [Step 25](#).

NOTE

In a normal cluster, ACL permission control is disabled for HBase by default. The **hbase:acl** table is generated only after ACL permission control is manually enabled. In this case, you need to check this table.

Step 19 View the HMaster startup status.

On the **Tasks** page shown in [Figure 7-103](#), the **RUNNING** value in the **State** column indicates that HMaster is being started and provides how much time HMaster keeps in that state. As shown in [Figure 7-104](#), if the state is **COMPLETE**, HMaster has been started.

Check whether HMaster has been in the **RUNNING** state for a long time.

Figure 7-103 HMaster being started

| Start Time | Description | State | Status |
|------------------------------|----------------|--------------------------|-------------------------------------|
| Thu Jan 28 14:43:12 CST 2016 | Master startup | RUNNING (since 1sec ago) | Initializing master service threads |

Figure 7-104 HMaster startup completed

| Start Time | Description | State | Status |
|------------------------------|----------------|----------------------------|--|
| Thu Jan 28 14:33:24 CST 2016 | Master startup | COMPLETE (since 59sec ago) | Calling postStartMaster coprocessors (since 56sec ago) |

- If yes, go to [Step 20](#).
- If no, go to [Step 21](#).

Step 20 On the HMaster web UI, check whether any **hbase:meta** is in the **Regions in Transition** state for a long time.

Figure 7-105 Regions in Transition

| Region | State | RIT time (ms) |
|--|--|---------------|
| 1588230740 | hbase.meta, 1588230740 state=PENDING_OPEN, ts=Wed Jan 27 19:49:27 CST 2016 (0s ago), server=10-64-35-147.21302.1453684877597 | 952 |
| Total number of Regions in Transition for more than 60000 milliseconds | | 0 |
| Total number of Regions in Transition | | 1 |

- If yes, go to [Step 21](#).
- If no, go to [Step 22](#).

Step 21 After ensuring that services are not affected, log in to FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services > HBase**, click **More**, and select **Restart Service**. In the dialog box that is displayed, enter the password, and click **OK**.

- If yes, go to [Step 22](#).
- If no, go to [Step 25](#).

NOTICE

During HBase service restart, the service is unavailable. For example, data cannot be read or written, table operations cannot be performed, and the HBase web UI is inaccessible.

Step 22 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 25](#).

Check whether the HBase configurations are correctly modified.

Step 23 On FusionInsight Manager, choose **Audit**. On the **Audit** page, click **Advanced Search**, click **...** on the right of **Operation Type**, select **Save configuration**, click **OK**, and click **Search**.

Step 24 In the search result, check whether the historical configurations of HBase-related services in the **Service** column, such as ZooKeeper, HDFS, and HBase, may affect the HBase service status. [Table 7-105](#) lists some configurations that may affect the HBase service status.

Table 7-105 Configurations affecting the HBase service status

| Parameter | Possible Impact |
|-----------|--|
| GC_OPTS | The memory configuration may be improper. You need to check the health status of instance processes. |

| Parameter | Possible Impact |
|--|---|
| hbase.rpc.protection | If the HBase service is not restarted offline after the value of this parameter is changed, the connection authentication fails and the HBase service becomes abnormal. |
| hbase.regionserver.metahandler.count | If there are too many regions in the cluster but this parameter is set to a small value, RIT may occur and regions cannot be brought online for a long time. |
| hbase.regionserver.thread.compaction.large | If this parameter is set to a large value, the node CPU usage may be too high. |
| hbase.regionserver.thread.compaction.small | If this parameter is set to a large value, the node CPU usage may be too high. |
| hbase.coprocessor.master.classes | If a custom coprocessor is used in the configuration, a logic error may cause the service to be unavailable. |
| hbase.coprocessor.region.classes | If a custom coprocessor is used in the configuration, a logic error may cause the service to be unavailable. |
| hbase.coprocessor.regionserver.classes | If a custom coprocessor is used in the configuration, a logic error may cause the service to be unavailable. |
| zookeeper.session.timeout | If this parameter is set to a small value, the connection between HBase and ZooKeeper times out too quickly. As a result, the HMaster instance and RegionServer may restart repeatedly. |

Check the network connection between HMaster and dependent components.

- Step 25** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > HBase**.
- Step 26** Click **Instances**. In the HMaster instance list, record the management IP address of the active HMaster instance.
- Step 27** Log in to the active HMaster node as user **omm** through the IP address obtained in [Step 26](#).
- Step 28** Run the **ping** command to check whether the network connection between the active HMaster node and the host where the dependent components reside is normal. (The dependent components include ZooKeeper, HDFS, and Yarn. The method of obtaining the IP address of the host where the dependent components reside is the same as that of obtaining the IP address of the active HMaster node.)
- If yes, go to [Step 31](#).
 - If no, go to [Step 29](#).

Step 29 Contact the network administrator to restore the network.

Step 30 In the alarm list, check whether this alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 31](#).

Collect fault information.

Step 31 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 32 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select the following services for the target cluster:

- ZooKeeper
- HDFS
- HBase

Step 33 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 34 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.168 ALM-19006 HBase Replication Sync Failed

Description

The alarm module checks the HBase DR data synchronization status every 30 seconds. When disaster recovery (DR) data fails to be synchronized to a standby cluster, the alarm is triggered.

When DR data synchronization succeeds, the alarm is cleared.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 19006 | Critical | Yes |

Parameters

| Name | Meaning |
|-------------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated. |

Impact on the System

HBase data in the cluster cannot be synchronized to the standby cluster. Synchronization data is stacked, causing a large amount of active/standby data inconsistency. As a result, the latest data cannot be read from the standby cluster after an active/standby DR switchover or dual-read. If the alarm persists, the storage space of the primary cluster and ZooKeeper nodes will be stacked, leading to service faults in the primary cluster.

Possible Causes

- The HBase service on the standby cluster is abnormal.
- A network exception occurs.

Procedure

Observe whether the system automatically clears the alarm.

- Step 1** On the FusionInsight Manager portal of the active cluster, click **O&M > Alarm > Alarms**.
- Step 2** In the alarm list, click the alarm to obtain alarm generation time from **Generated** of the alarm. Check whether the alarm has existed for five minutes.
- If yes, go to **Step 4**.
 - If no, go to **Step 3**.
- Step 3** Wait five minutes and check whether the system automatically clears the alarm.
- If yes, no further action is required.
 - If no, go to **Step 4**.

Check the HBase service status of the standby cluster.

Step 4 Log in to the FusionInsight Manager portal of the active cluster, and click **O&M > Alarm > Alarms**.

Step 5 In the alarm list, click the alarm to obtain **HostName** from **Location**.

Step 6 Access the node where the HBase client of the active cluster resides as user **omm**.

If the cluster uses a security mode, perform security authentication first and then access the **hbase shell** interface as user **hbase**.

```
cd /opt/client
```

```
source ./bigdata_env
```

```
kinit hbaseuser
```

Step 7 Run the **status 'replication', 'source'** command to check the DR synchronization status of the faulty node.

The DR synchronization status of a node is as follows.

```
10-10-10-153:
SOURCE: PeerID=abc, SizeOfLogQueue=0, ShippedBatches=2, ShippedOps=2, ShippedBytes=320,
LogReadInBytes=1636, LogEditsRead=5, LogEditsFiltered=3, SizeOfLogToReplicate=0,
TimeForLogToReplicate=0, ShippedHFiles=0, SizeOfHFileRefsQueue=0, AgeOfLastShippedOp=0,
TimeStampsOfLastShippedOp=Mon Jul 18 09:53:28 CST 2016, Replication Lag=0,
FailedReplicationAttempts=0
SOURCE: PeerID=abc1, SizeOfLogQueue=0, ShippedBatches=1, ShippedOps=1, ShippedBytes=160,
LogReadInBytes=1636, LogEditsRead=5, LogEditsFiltered=3, SizeOfLogToReplicate=0,
TimeForLogToReplicate=0, ShippedHFiles=0, SizeOfHFileRefsQueue=0, AgeOfLastShippedOp=16788,
TimeStampsOfLastShippedOp=Sat Jul 16 13:19:00 CST 2016, Replication Lag=16788,
FailedReplicationAttempts=5
```

Step 8 Obtain **PeerID** corresponding to a record whose **FailedReplicationAttempts** value is greater than 0.

In the preceding step, data on the faulty node 10-10-10-153 fails to be synchronized to a standby cluster whose **PeerID** is **abc1**.

Step 9 Run the **list_peers** command to find the cluster and the HBase instance corresponding to the **PeerID** value.

```
PEER_ID CLUSTER_KEY STATE TABLE_CFS
abc1 10.10.10.110,10.10.10.119,10.10.10.133:2181:/hbase2 ENABLED
abc 10.10.10.110,10.10.10.119,10.10.10.133:2181:/hbase ENABLED
```

In the preceding information, **/hbase2** indicates that data is synchronized to the HBase2 instance of the standby cluster.

Step 10 In the service list of FusionInsight Manager of the standby cluster, check whether the running status of the HBase instance obtained by using **Step 9** is **Normal**.

- If yes, go to **Step 14**.
- If no, go to **Step 11**.

Step 11 In the alarm list, check whether the **ALM-19000 HBase Service Unavailable** alarm is generated.

- If yes, go to **Step 12**.
- If no, go to **Step 14**.

Step 12 Follow troubleshooting procedures in **ALM-19000 HBase Service Unavailable** to rectify the fault.

Step 13 Wait for a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 14](#).

Check network connections between RegionServers on active and standby clusters.

Step 14 Log in to the FusionInsight Manager portal of the active cluster, and click **O&M > Alarm > Alarms**.

Step 15 In the alarm list, click the alarm to obtain **HostName** from **Location**.

Step 16 Use the IP address obtained in [Step 15](#) to log in to a faulty RegionServer node as user **omm**.

Step 17 Run the **ping** command to check whether network connections between the faulty RegionServer node and the host where RegionServer of the standby cluster resides are in the normal state.

- If yes, go to [Step 20](#).
- If no, go to [Step 18](#).

Step 18 Contact the network administrator to restore the network.


Step 19 After the network is running properly, check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to [Step 20](#).

Collect fault information.

Step 20 On the FusionInsight Manager interface of active and standby clusters, choose **O&M > Log > Download**.

Step 21 In the **Service** drop-down list box, select **HBase** in the required cluster.

Step 22 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 23 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.169 ALM-19007 HBase GC Time Exceeds the Threshold

Description

The system checks the old generation garbage collection (GC) time of the HBase service every 60 seconds. This alarm is generated when the detected old generation GC time exceeds the threshold (exceeds 5 seconds for three consecutive checks by default). To change the threshold, on the FusionInsight Manager portal, choose **O&M > Alarm > Thresholds > Name of the desired cluster > HBase > GC > GC time for old generation**. This alarm is cleared when the old generation GC time of the HBase service is shorter than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 19007 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |

Impact on the System

If the GC time of the old generation exceeds the threshold, the read and write of HBase data will slow down. In severe cases, the request times out.

Possible Causes

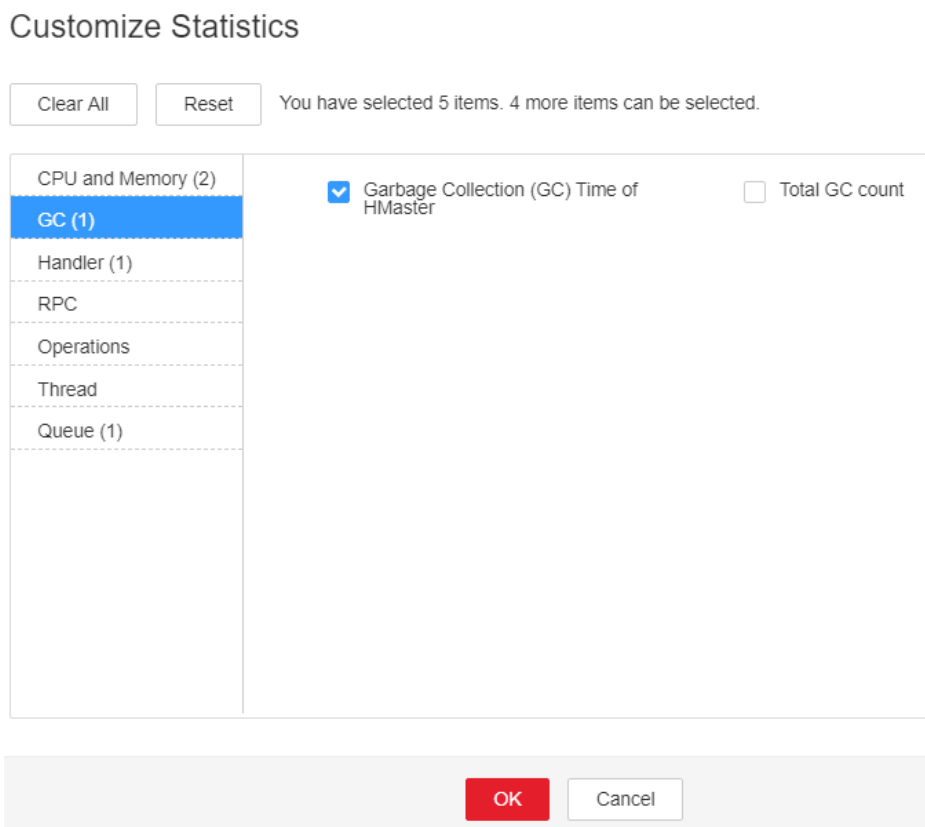
The memory of HBase instances is overused, the heap memory is inappropriately allocated, or a large number of I/O operations exist in HBase. As a result, GCs occur frequently.

Procedure

Check the GC time.

- Step 1** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **ID** is **19007**. Then check the role name in **Location** and confirm the IP address of the instance.
- If the role for which the alarm is generated is HMaster, go to [Step 2](#).
 - If the role for which the alarm is generated is RegionServer, go to [Step 3](#).
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HBase > Instance** and click the HMaster for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize > GC > Garbage Collection (GC) Time of HMaster** and click **OK** to check whether the value of **GC time for old generation** is greater than the threshold (exceeds 5 seconds for three consecutive checks periods by default).
- If yes, go to [Step 4](#).
 - If no, go to [Step 6](#).

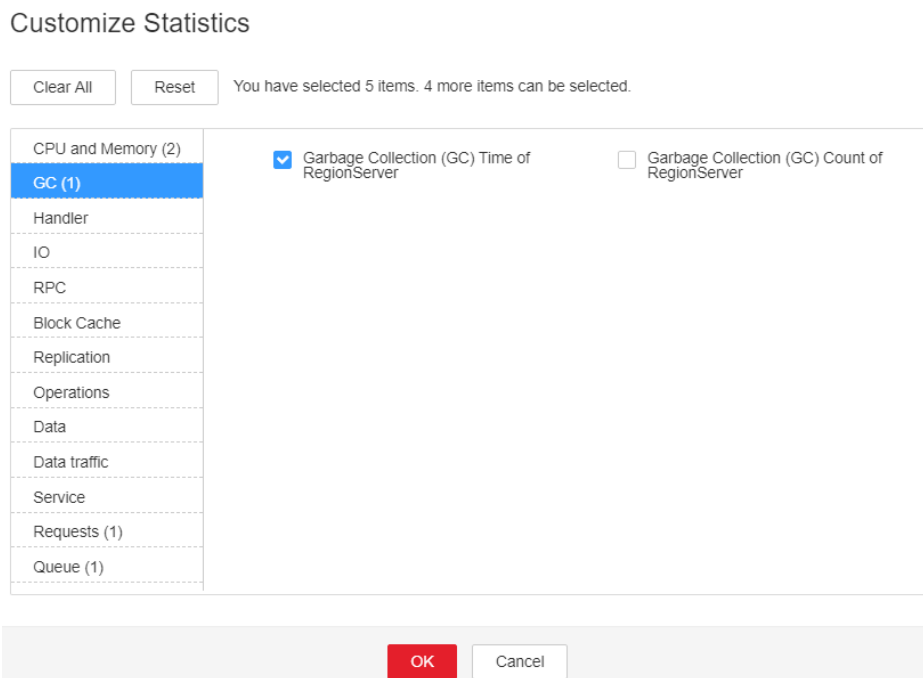
Figure 7-106 Garbage Collection (GC) Time of HMaster



- Step 3** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HBase > Instance** and click the RegionServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize > GC > Garbage Collection (GC) Time of RegionServer** and click **OK** to check whether the value of **GC time for old generation** is greater than the threshold (exceeds 5 seconds for three consecutive checks periods by default).
- If yes, go to [Step 4](#).

- If no, go to [Step 6](#).

Figure 7-107 Garbage Collection (GC) Time of RegionServer



Check the current JVM configuration.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations**, and click **All Configurations**. In Search, enter **GC_OPTS** to check the **GC_OPTS** memory parameter of role HMaster(HBase->HMaster), RegionServer(HBase->RegionServer). Adjust the values of **-Xmx** and **-XX:CMSInitiatingOccupancyFraction** of the GC_OPTS parameter by referring to the Note.

 NOTE

1. Suggestions on GC parameter configurations for HMaster
 - Set **-Xms** and **-Xmx** to the same value to prevent JVM from dynamically adjusting the heap memory size and affecting performance.
 - Set **-XX:NewSize** to the value of **-XX:MaxNewSize**, which is one eighth of **-Xmx**.
 - For large-scale HBase clusters with a large number of regions, increase values of **GC_OPTS** parameters for HMaster. Specifically, set **-Xmx** to 4 GB if the number of regions is less than 100,000. If the number of regions is more than 100,000, set **-Xmx** to be greater than or equal to 6 GB. For each increased 35,000 regions, increase the value of **-Xmx** by 2 GB. The maximum value of **-Xmx** is 32 GB.
2. Suggestions on GC parameter configurations for RegionServer
 - Set **-Xms** and **-Xmx** to the same value to prevent JVM from dynamically adjusting the heap memory size and affecting performance.
 - Set **-XX:NewSize** to one eighth of **-Xmx**.
 - Set the memory for RegionServer to be greater than that for HMaster. If sufficient memory is available, increase the heap memory.
 - Set **-Xmx** based on the machine memory size. Specifically, set **-Xmx** to 32 GB if the machine memory is greater than 200 GB, to 16 GB if the machine memory is greater than 128 GB and less than 200 GB, and to 8 GB if the machine memory is less than 128 GB. When **-Xmx** is set to 32 GB, a RegionServer node supports 2000 regions and 200 hotspot regions.
 - **XX:CMSInitiatingOccupancyFraction** to be less than and equal to **85**, and it is calculated as follows: $100 \times (\text{hfile.block.cache.size} + \text{hbase.regionserver.global.memstore.size})$


Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager interface of active and standby clusters, choose **O&M > Log > Download**.

Step 7 In the **Service** drop-down list box, select **HBase** in the required cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.170 ALM-19008 Heap Memory Usage of the HBase Process Exceeds the Threshold

Description

The system checks the HBase service status every 30 seconds. The alarm is generated when the heap memory usage of an HBase service exceeds the threshold (90% of the maximum memory).

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 19008 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |

Impact on the System

The available HBase memory is insufficient, which may cause node restart. During the node restart, the read/write request delay on the node increases or fails.

Possible Causes

The heap memory of the HBase service is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

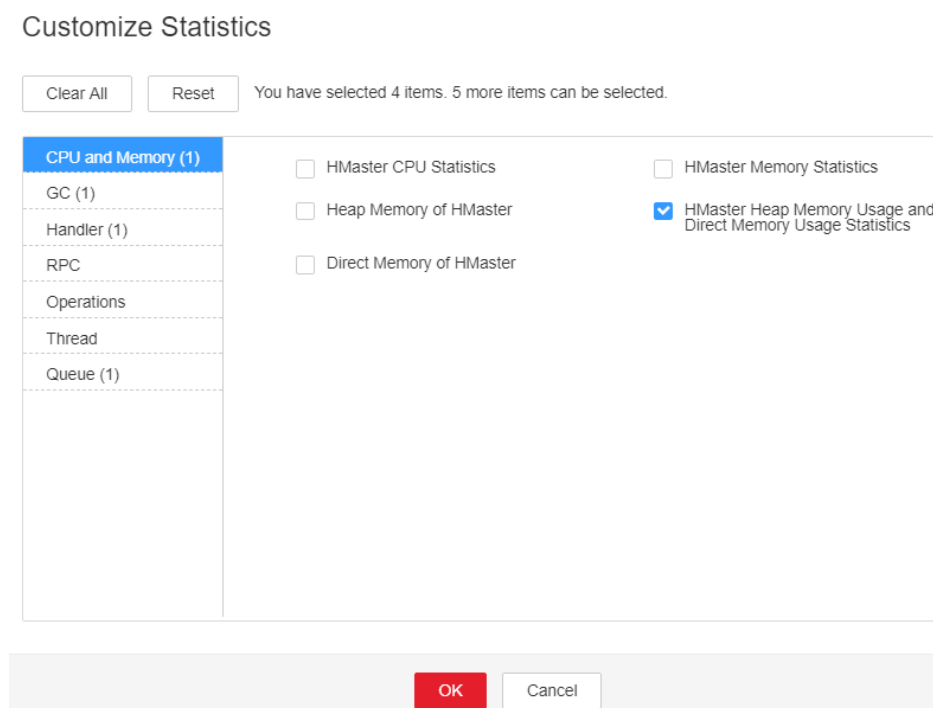
- Step 1** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **ID** is **19008**. Then check the role name in **Location** and confirm the IP address of the instance.
 - If the role for which the alarm is generated is HMaster, go to [Step 2](#).

- If the role for which the alarm is generated is RegionServer, go to [Step 3](#).

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the HMaster for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory** > **HMaster Heap Memory Usage and Direct Memory Usage Statistics** and click **OK**, check whether the used heap memory of the HBase service reaches 90% of the maximum heap memory specified for HBase.

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

Figure 7-108 HMaster Heap Memory Usage and Direct Memory Usage Statistics



Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the RegionServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory** > **RegionServer Heap Memory Usage and Direct Memory Usage Statistics** and click **OK**, check whether the used heap memory of the HBase service reaches 90% of the maximum heap memory specified for HBase.

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

Figure 7-109 RegionServer Heap Memory Usage and Direct Memory Usage Statistics

Customize Statistics

You have selected 4 items. 5 more items can be selected.

| | | |
|---------------------------|--|---|
| CPU and Memory (1) | <input type="checkbox"/> RegionServer CPU Statistics | <input type="checkbox"/> RegionServer Memory Statistics |
| GC (1) | <input type="checkbox"/> Heap Memory of RegionServer | <input checked="" type="checkbox"/> RegionServer Heap Memory Usage and Direct Memory Usage Statistics |
| Handler | | |
| IO | <input type="checkbox"/> Direct Memory of RegionServer | |
| RPC | | |
| Block Cache | | |
| Replication | | |
| Operations | | |
| Data | | |
| Data traffic | | |
| Service | | |
| Requests (1) | | |
| Queue (1) | | |

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations**, and click **All Configurations**. Choose **HMaster/RegionServer** > **System**. Increase the value of **-Xmx** in **GC_OPTS** by referring to the Note.

 **NOTE**

1. Suggestions on GC parameter configurations for HMaster
 - Set **-Xms** and **-Xmx** to the same value to prevent JVM from dynamically adjusting the heap memory size and affecting performance.
 - Set **-XX:NewSize** to the value of **-XX:MaxNewSize**, which is one eighth of **-Xmx**.
 - For large-scale HBase clusters with a large number of regions, increase values of **GC_OPTS** parameters for HMaster. Specifically, set **-Xmx** to 4 GB if the number of regions is less than 100,000. If the number of regions is more than 100,000, set **-Xmx** to be greater than or equal to 6 GB. For each increased 35,000 regions, increase the value of **-Xmx** by 2 GB. The maximum value of **-Xmx** is 32 GB.
2. Suggestions on GC parameter configurations for RegionServer
 - Set **-Xms** and **-Xmx** to the same value to prevent JVM from dynamically adjusting the heap memory size and affecting performance.
 - Set **-XX:NewSize** to one eighth of **-Xmx**.
 - Set the memory for RegionServer to be greater than that for HMaster. If sufficient memory is available, increase the heap memory.
 - Set **-Xmx** based on the machine memory size. Specifically, set **-Xmx** to 32 GB if the machine memory is greater than 200 GB, to 16 GB if the machine memory is greater than 128 GB and less than 200 GB, and to 8 GB if the machine memory is less than 128 GB. When **-Xmx** is set to 32 GB, a RegionServer node supports 2000 regions and 200 hotspot regions.


Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M > Log > Download**.

Step 7 Select **HBase** in the required cluster from the **Service** drop-down list.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.171 ALM-19009 Direct Memory Usage of the HBase Process Exceeds the Threshold

Description

The system checks the HBase service status every 30 seconds. The alarm is generated when the direct memory usage of an HBase service exceeds the threshold (90% of the maximum memory).

The alarm is cleared when the direct memory usage is less than the threshold.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 19009 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|--|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service name for which the alarm is generated. |

| Name | Meaning |
|----------|--|
| RoleName | Specifies the role name for which the alarm is generated. |
| HostName | Specifies the object (host ID) for which the alarm is generated. |

Impact on the System

The available HBase direct memory is insufficient, which may cause node restart. During the node restart, the read/write request delay on the node increases or fails.

Possible Causes

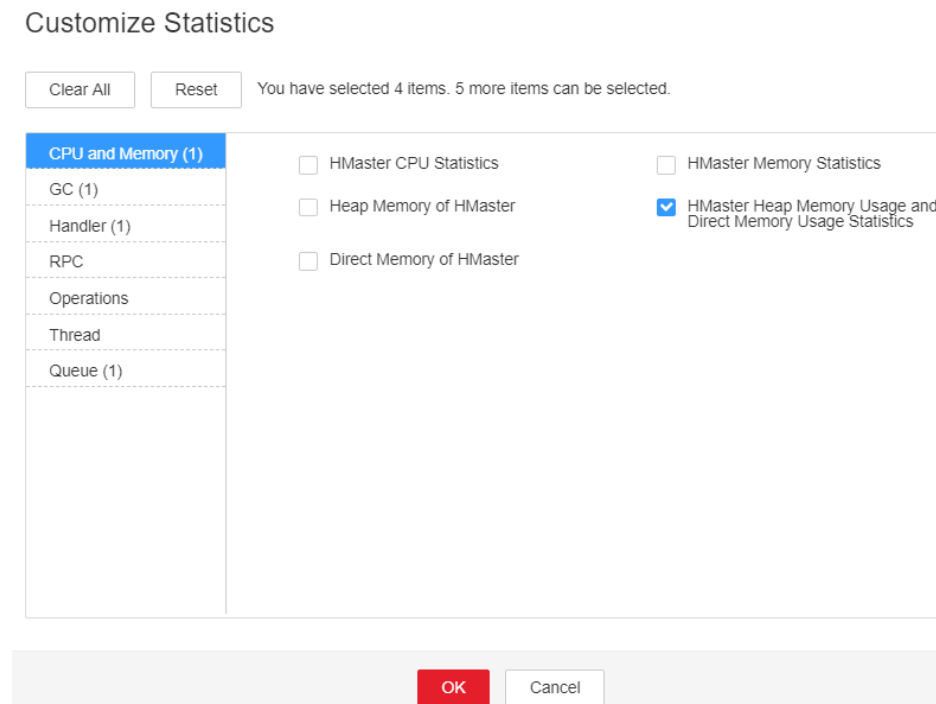
The direct memory of the HBase service is overused or the direct memory is inappropriately allocated.

Procedure


Check direct memory usage.

- Step 1** On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **ID** is **19009**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- If the role for which the alarm is generated is HMaster, go to [Step 2](#).
 - If the role for which the alarm is generated is RegionServer, go to [Step 3](#).
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > HBase > Instance** and click the HMaster for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize > CPU and Memory > HMaster Heap Memory Usage and Direct Memory Usage Statistics** and click **OK** to check whether the used direct memory of the HBase service reaches 90% of the maximum direct memory specified for HBase.
- If yes, go to [Step 4](#).
 - If no, go to [Step 8](#).

Figure 7-110 HMaster Heap Memory Usage and Direct Memory Usage Statistics



- Step 3** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the RegionServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory** > **RegionServer Heap Memory Usage and Direct Memory Usage Statistics** and click **OK** to check whether the used direct memory of the HBase service reaches 90% of the maximum direct memory specified for HBase.
- If yes, go to **Step 4**.
 - If no, go to **Step 8**.

Figure 7-111 RegionServer Heap Memory Usage and Direct Memory Usage Statistics

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations**, and click **All Configurations**. Choose **HMaster/RegionServer** > **System** and check whether **XX:MaxDirectMemorySize** exists in **GC_OPTS**.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations**, and click **All Configurations**. Choose **HMaster/RegionServer** > **System** and delete **XX:MaxDirectMemorySize** from **GC_OPTS**.

Step 6 Check whether the **ALM-19008 Heap Memory Usage of the HBase Process Exceeds the Threshold** alarm is generated.

If yes, handle the alarm by referring to **ALM-19008 Heap Memory Usage of the HBase Process Exceeds the Threshold**.

If no, go to **Step 8**.


Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On the FusionInsight Manager interface of active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 9 In the **Service** in the required cluster drop-down list box, select **HBase**.

Step 10 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.172 ALM-19011 RegionServer Region Number Exceeds the Threshold

Description

The system checks the number of regions on each RegionServer in each HBase service instance every 30 seconds. The region number is displayed on the HBase service monitoring page and RegionServer role monitoring page. This alarm is generated when the number of regions on a RegionServer exceeds the threshold (default value: 2000) for 20 consecutive times. The threshold can be changed by choosing **O&M > Alarm > Thresholds > Name of the desired cluster > HBase**. This alarm is cleared when the number of regions is less than or equal to the threshold.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 19011 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

If the number of RegionServer regions exceeds the threshold, too many Regions increase the load of RegionServer, causing resource bottlenecks such as memory, disk I/O, and CPU. As a result, request response becomes slow or even times out.

Possible Causes

- The RegionServer region distribution is unbalanced.
- The HBase cluster scale is too small.

Procedure

View alarm location information.

Step 1 On the FusionInsight Manager home page, choose **O&M > Alarm > Alarms**, select this alarm, and view the service instance and host name in **Location**.


Step 2 On the FusionInsight Manager home page, choose **Cluster > Name of the desired cluster > Services**, click the HBase service instance for which the alarm is generated, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, check whether the region distribution on the RegionServer is balanced.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

- If yes, go to [Step 9](#).
- If no, go to [Step 3](#).

Figure 7-112 WebUI of HBase instance



| ServerName | Start time | Last contact | Version | Requests Per Second | Num. Regions |
|-------------|------------------------------|--------------|---------|---------------------|--------------|
| 100-0173571 | Fri Sep 11 15:42:53 CST 2020 | 1 s | 2.2.3 | 0 | 8 |
| 100-0173988 | Fri Sep 11 15:42:53 CST 2020 | 0 s | 2.2.3 | 0 | 4 |
| 100-0173989 | Fri Sep 11 15:42:53 CST 2020 | 1 s | 2.2.3 | 0 | 4 |
| Total: | | | | 0 | 16 |

Enable load balancing.

Step 3 Log in to the node where the HBase client is located as user **root**. Go to the client installation directory, and set environment variables.

```
cd client installation directory
```

```
source bigdata_env
```

If the cluster adopts the security mode, perform security authentication. Specifically, run the **kinit hbase** command and enter the password as prompted (obtain the password from the administrator).

Step 4 Run the following commands to go to the HBase shell command window and check whether the load balancing function is enabled.

```
hbase shell
```

balancer_enabled

- If yes, go to [Step 6](#).
- If no, go to [Step 5](#).

Step 5 On the HBase shell command window, run the following commands to enable the load balancing function and check whether the function is enabled.

balance_switch true

balancer_enabled

Step 6 On the HBase shell command window, run the **balancer** command to manually trigger the load balancing function.

 **NOTE**

You are advised to enable and manually trigger the load balancing function during off-peak hours.

Step 7 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, refresh the page and check whether the region distribution is balanced.

- If yes, go to [Step 8](#).
- If no, go to [Step 21](#).

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Delete unwanted HBase tables.

 **NOTE**

Exercise caution when deleting data to ensure data is deleted correctly.

Step 9 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, view tables stored in the HBase service instance and record unwanted tables that can be deleted.

Step 10 On the HBase shell command window, run the **disable** command and **drop** command to delete the table to decrease the number of regions.

disable '*name of the table to be deleted*'

drop '*name of the table to be deleted*'

Step 11 On the HBase shell command window, run the following command to check whether the load balancing function is enabled.

balancer_enabled

- If yes, go to [Step 13](#).
- If no, go to [Step 12](#).

Step 12 On the HBase shell command window, run the following commands to enable the load balancing function and confirm that the function is enabled.

balance_switch true

balancer_enabled

Step 13 On the HBase shell command window, run the **balancer** command to manually trigger the load balancing function.

Step 14 On the FusionInsight Manager home page, choose **Cluster > Name of the desired cluster > Services > HBase**, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, refresh the page and check whether the region distribution is balanced.

- If yes, go to [Step 15](#).
- If no, go to [Step 21](#).

Step 15 Check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 16](#).

Adjust the threshold.

Step 16 On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Name of the desired cluster > HBase > Regions(RegionServer)**, select the applied rule, and click **Modify** to check whether the threshold is proper.

- If it is excessively small, increase the threshold as required and go to [Step 17](#).
- If it is proper, go to [Step 18](#).

Figure 7-113 Regions(RegionServer_1)



Step 17 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 18](#).

Perform system capacity expansion.

Step 18 Add nodes to the HBase cluster and add RegionServer instances to the nodes. Then enable and manually trigger the load balancing function.

Step 19 On the FusionInsight Manager home page, choose **Cluster > Name of the desired cluster > Services**, click the HBase service instance for which the alarm is generated, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, refresh the page and check whether the region distribution is balanced.

- If yes, go to [Step 20](#).
- If no, go to [Step 21](#).

Step 20 Check whether the alarm is cleared.


- If yes, no further action is required.

- If no, go to [Step 21](#).

Collect fault information.

Step 21 On the FusionInsight Manager home page of the active and standby clusters, choose **O&M> Log > Download**.

Step 22 Select **HBase** in the required cluster from the **Service**.

Step 23 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 24 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.173 ALM-19012 HBase System Table Directory or File Lost

Description

The system checks whether HBase directories and files exist on the HDFS every 120 seconds. This alarm is generated when the system detects that the files or directories do not exist. This alarm is cleared when the files or directories are restored.

The HBase directories and files are as follows:

- Directory of the namespace **hbase** on the HDFS
- **hbase.version** file
- Directory of the table **hbase:meta** on the HDFS, .tableinfo file, and .regioninfo file
- Directory of the table **hbase:namespace** on the HDFS, .tableinfo file, and .regioninfo file
- Directory of the table **hbase:hindex** on the HDFS, .tableinfo file, and .regioninfo file
- Directory of the **hbase:acl** table on the HDFS, .tableinfo, and .regioninfo file (This table does not exist in the common mode cluster by default.)

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 19012 | Critical | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System


The HBase service fails to be restarted or started. As a result, all HBase service requests fail.

Possible Causes

Files or directories on the HDFS are missing.

Procedure

Locate the alarm cause.

- Step 1** On the FusionInsight Manager, choose **O&M > Alarm > Alarms**. Click this alarm and check whether **Alarm Cause** indicates unknown errors.
- If yes, go to **Step 4**.
 - If no, go to **Step 2**
- Step 2** On the FusionInsight Manager home page, choose **O&M > Backup and Restoration > Backup Management**. Check whether there are success records of the backup task named **default** or other HBase metadata backup tasks that have been successfully executed.
- If yes, go to **Step 3**.
 - If no, go to **Step 4**.
- Step 3** Use the latest backup metadata to restore the metadata of the HBase service.
- Collect fault information.**
- Step 4** On the FusionInsight Manager page of the active and standby clusters, choose **O&M > Log > Download**.
- Step 5** In the **Service** area, select faulty HBase services in the required cluster.
- Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.174 ALM-19013 Duration of Regions in transaction State Exceeds the Threshold

Description

The system checks the number of regions in transaction state on HBase every 300 seconds. This alarm is generated when the system detects that the duration of regions in transaction state exceeds the threshold for two consecutive times. This alarm is cleared when all timeout regions are restored.

Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 19013 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

Some data in the table gets lost or becomes unavailable.

Possible Causes

- Compaction is permanently blocked.
- The HDFS files are abnormal.

Procedure

Locate the alarm cause.

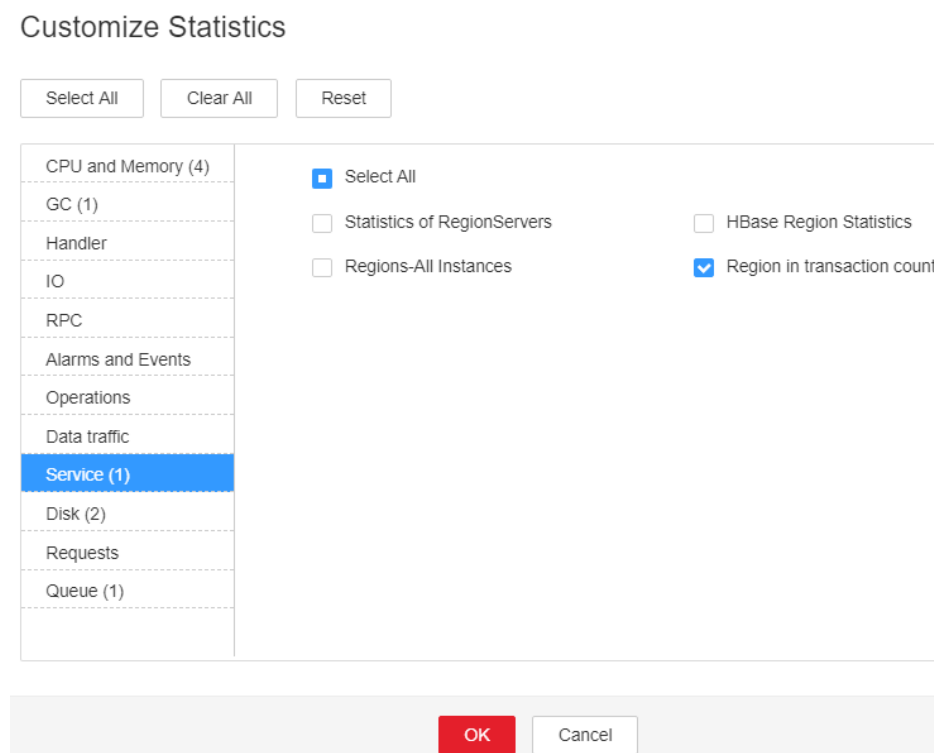
Step 1 On the FusionInsight Manager, choose **O&M > Alarm > Alarms**, select this alarm, and view the **HostName** and **RoleName** in **Location**.

Step 2 Choose **Cluster > Name of the desired cluster > Services > HBase**, Click the drop-down menu in the chartarea and choose **Customize > Service >**

Region in transaction count to view **Region in transaction count over threshold**. Check whether the monitoring item detects a value in three consecutive detection periods. (The default threshold is 60 seconds.)

- If yes, go to [Step 3](#).
- If no, go to [Step 7](#).

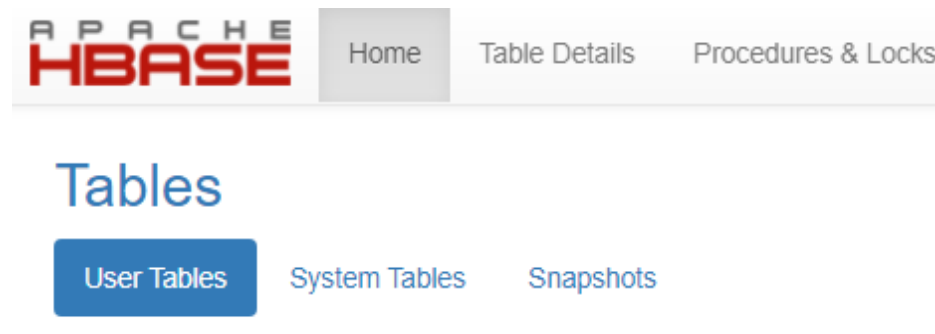
Figure 7-114 Region in transaction count



Step 3 Choose **Cluster > Name of the desired cluster > Services > HBase > HMaster (Active) > Tables** to check whether the regions of only one table transaction status time out.

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

Figure 7-115 Tables



Step 4 Run the **hbase hbck** command on the client and check whether the error message "No table descriptor file under hdfs://hacluster/hbase/data/default/table" is displayed.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Log in to the client as user **root**. Run the following command:

```
cd client installation directory
```

```
source bigdata_env
```

If the cluster is in security mode, run the **kinit hbase** command

Log in to the HMaster WebUI, choose **Procedure & Locks** in the navigation tree, and check whether any process ID is in the **Waiting** state in **Procedures**. If yes, run the following command to release the procedure lock:

```
hbase hbck -j client installation directory/HBase/hbase/tools/hbase-hbck2-*.jar  
bypass -o pid
```

Check whether the state is in the **Bypass** state. If the procedure on the UI is always in **RUNNABLE(Bypass)** state, perform an active/standby switchover. Run the **assigns** command to bring the region online again.

```
hbase hbck -j client installation directory/HBase/hbase/tools/hbase-hbck2-*.jar  
assigns -o regionName
```


Step 6 Repeat **Step 4**. Run the **hbase hbck** command on the client and check whether the error message "No table descriptor file under hdfs://hacluster/hbase/data/default/table" is displayed.

- If yes, go to **Step 7**.
- If no, no further action is required.

Collect fault information.

Step 7 On the FusionInsight Manager page of the active and standby clusters, choose **O&M > Log > Download**.

Step 8 In the **Service** area, select faulty HBase services in the required cluster.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

7.12.175 ALM-19014 Capacity Quota Usage on ZooKeeper Exceeds the Threshold Severely

Alarm Description

The system checks the ZNode usage of the HBase service every 120 seconds. This alarm is generated when the ZNode capacity usage of the HBase service exceeds the critical alarm threshold (90% by default).

This alarm is cleared when the ZNode capacity usage is less than the critical alarm threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19014 | Critical | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Threshold | Specifies the threshold for generating the alarm. |

Impact on the System

This alarm indicates that the capacity usage of the ZNode of HBase has exceeded the threshold severely. As a result, the write request of the HBase service fails.

Possible Causes

- DR is configured for HBase, and data synchronization fails or is slow in DR.
- A large number of WAL files are being split in the HBase cluster.

Handling Procedure

Check the capacity configuration and usage of ZNodes.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**, select the alarm whose ID is **19014**, and view the threshold in **Additional Information**.

Step 2 Log in to the HBase client as user **root**. Run the following command to go to the client installation directory:

```
cd Client installation directory
```

Run the following command to set environment variables:

```
source bigdata_env
```

If the cluster uses the security mode, run the following command to perform security authentication:

```
kinit hbase
```

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 3 Run the **hbase zkcli** command to log in to the ZooKeeper client and run the **listquota /hbase** command to check the ZNode capacity quota of the HBase service. The ZNode root directory in the command is specified by the **zookeeper.znode.parent** parameter of the HBase service. The marked area in the following figure shows the capacity configuration of the root ZNode of the HBase service.


```
[zk: :24002, :24002, :24002(CONNECTED) 145] listquota /hbase
absolute path is /zookeeper/quota/hbase
Output quota for /hbase count=1500000,bytes=10240
Output stat for /hbase count=42,bytes=1601
```

Step 4 Run the **getusage /hbase/splitWAL** command to check the capacity usage of the ZNode. Check whether the ratio of **Data size** to the ZNode capacity quota is close to the alarm threshold.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Check whether the alarm whose ID is **12007**, **19000**, or **19013** and the **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.

- If no, go to [Step 9](#).
- Step 6** Run the `getusage /hbase/replication` command to check the capacity usage of the ZNode. Check whether the ratio of **Data size** to the ZNode capacity quota is close to the alarm threshold.
- If yes, go to [Step 7](#).
 - If no, go to [Step 9](#).
- Step 7** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Check whether the alarm whose ID is **19006** and **ServiceName** in **Location** is the current HBase service exists.
- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to [Step 8](#).
 - If no, go to [Step 9](#).
- Step 8** Check whether the alarm is cleared five minutes later.
- If yes, no further action is required.
 - If no, go to [Step 9](#).
- Collect the fault information.**
- Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 10** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HBase** for the target cluster.
- Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 12** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.176 ALM-19015 Quantity Quota Usage on ZooKeeper Exceeds the Threshold

Alarm Description

The system checks the ZNode usage of the HBase service every 120 seconds. This alarm is generated when the system detects that the ZNode quantity usage of the HBase service exceeds the alarm threshold (75% by default).

This alarm is cleared when the ZNode quantity usage is less than the alarm threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19015 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Threshold | Specifies the threshold for generating the alarm. |

Impact on the System

This alarm indicates that the ZNode quantity usage in the HBase service has exceeded the threshold. If this alarm is not handled in a timely manner, the problem severity may be escalated to **Critical** and data fails to be written.

Possible Causes

- DR is configured for HBase, and data synchronization fails or is slow in DR.
- A large number of WAL files are being split in the HBase cluster.

Handling Procedure

Check the quantity quota and usage of ZNodes.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**, select the alarm whose ID is **19015**, and view the threshold in **Additional Information**.

Step 2 Log in to the HBase client as user **root**. Run the following command to go to the client installation directory:

```
cd Client installation directory
```

Run the following command to set environment variables:

```
source bigdata_env
```

If the cluster uses the security mode, run the following command to perform security authentication:

kinit hbase

Enter the password as prompted (obtain the password from the MRS cluster administrator).

- Step 3** Run the **hbase zkcli** command to log in to the ZooKeeper client and run the **listquota /hbase** command to check the ZNode quantity quota of the HBase service. The ZNode root directory in the command is specified by the **zookeeper.znode.parent** parameter of the HBase service. The marked area in the following figure shows the quantity quota configuration of the root ZNode of the HBase service.

```
[zk: :24002, :24002, :24002(CONNECTED) 7] listquota /hbase
absolute path is /zookeeper/quota/hbase
Output quota for /hbase count=1500000,bytes=10240
Output stat for /hbase count=59,bytes=1902
```

- Step 4** Run the **getusage /hbase/splitWAL** command to check the ZNode quantity usage and check whether the ratio of **Node count** in the command output to the ZNode quantity quota is close to the alarm threshold.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

- Step 5** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Check whether the alarm whose ID is **12007**, **19000**, or **19013** and the **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to [Step 8](#).
- If no, go to [Step 9](#).

- Step 6** Run the **getusage /hbase/replication** command to check the ZNode quantity usage and check whether the ratio of **Node count** in the command output to the ZNode quantity quota is close to the alarm threshold.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

- Step 7** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Check whether the alarm whose ID is **19006** and **ServiceName** in **Location** is the current HBase service exists.


- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to [Step 8](#).
- If no, go to [Step 9](#).

- Step 8** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Collect the fault information.

- Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 10** Expand the **Service** drop-down list, and select **HBase** for the target cluster.
- Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 12** Contact O&M personnel and provide the collected logs.
- End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.177 ALM-19016 Quantity Quota Usage on ZooKeeper Exceeds the Threshold Severely

Alarm Description

The system checks the ZNode usage of the HBase service every 120 seconds. This alarm is generated when the znode usage of the HBase service exceeds the critical alarm threshold (90% by default).

This alarm is cleared when the quantity usage of the ZNode is less than the critical alarm threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19016 | Critical | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |

| Parameter | Description |
|-----------|---|
| Threshold | Specifies the threshold for generating the alarm. |

Impact on the System

This alarm indicates that the quantity usage of the ZNode of HBase has exceeded the threshold severely. As a result, the write request of the HBase service fails.

Possible Causes

- DR is configured for HBase, and data synchronization fails or is slow in DR.
- A large number of WAL files are being split in the HBase cluster.

Handling Procedure

Check the quantity quota and usage of ZNodes.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**, select the alarm whose ID is **19016**, and view the threshold in **Additional Information**.

Step 2 Log in to the HBase client as user **root**. Run the following command to go to the client installation directory:

```
cd Client installation directory
```

Run the following command to set environment variables:

```
source bigdata_env
```

If the cluster uses the security mode, run the following command to perform security authentication:

```
kinit hbase
```

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 3 Run the **hbase zkcli** command to log in to the ZooKeeper client and run the **listquota /hbase** command to check the ZNode quantity quota of the HBase service. The ZNode root directory in the command is specified by the **zookeeper.znode.parent** parameter of the HBase service. The marked area in the following figure shows the quantity configuration of the root ZNode of the HBase service.

```
[zk: :24002, :24002, :24002(CONNECTED) 7] listquota /hbase
absolute path is /zookeeper/quota/hbase
Output quota for /hbase count=1500000,bytes=10240
Output stat for /hbase count=59,bytes=1902
```

Step 4 Run the **getusage /hbase/splitWAL** command to check the ZNode usage and check whether the ratio of **Node count** in the command output to the znode quantity quota is close to the alarm threshold.

- If yes, go to [Step 5](#).

- If no, go to [Step 6](#).

Step 5 On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Check whether the alarm whose ID is **12007**, **19000**, or **19013** and the **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to [Step 8](#).
- If no, go to [Step 9](#).

Step 6 Run the **getusage /hbase/replication** command to check the ZNode usage and check whether the ratio of **Node count** in the command output to the ZNode quantity quota is close to the alarm threshold.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

Step 7 On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Check whether the alarm whose ID is **19006** and **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to [Step 8](#).
- If no, go to [Step 9](#).


Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 10 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.178 ALM-19017 Capacity Quota Usage on ZooKeeper Exceeds the Threshold

Alarm Description

The system checks the ZNode usage of the HBase service every 120 seconds. This alarm is generated when the system detects that the ZNodes capacity usage of the HBase service exceeds the alarm threshold (75% by default).

This alarm is cleared when the capacity usage of the ZNode capacity is less than the threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19017 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Threshold | Specifies the threshold for generating the alarm. |

Impact on the System

This alarm indicates that the ZNodes capacity usage in the HBase service has exceeded the threshold. If this alarm is not handled in a timely manner, the problem severity may be escalated to **Critical**, affecting data writing.

Possible Causes

- DR is configured for HBase, and data synchronization fails or is slow in DR.
- A large number of WAL files are being split in the HBase cluster.

Handling Procedure

Check the capacity configuration and usage of ZNodes.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**, select the alarm whose ID is **19017**, and view the threshold in **Additional Information**.

Step 2 Log in to the HBase client as user **root**. Run the following command to go to the client installation directory:

```
cd Client installation directory
```

Run the following command to set environment variables:

```
source bigdata_env
```

If the cluster uses the security mode, run the following command to perform security authentication:

```
kinit hbase
```

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 3 Run the **hbase zkcli** command to log in to the ZooKeeper client and run the **listquota /hbase** command to check the ZNode quantity quota of the HBase service. The ZNode root directory in the command is specified by the **zookeeper.znode.parent** parameter of the HBase service. The marked area in the following figure shows the quantity configuration of the root ZNode of the HBase service.

```
[zk: :24002, :24002, :24002(CONNECTED) 145] listquota /hbase
absolute path is /zookeeper/quota/hbase
Output quota for /hbase count=1500000,bytes=10240
Output stat for /hbase count=42,bytes=1601
```

Step 4 Run the **getusage /hbase/splitWAL** command to check the capacity usage of the ZNode. Check whether the ratio of **Data size** to the ZNode capacity quota is close to the alarm threshold.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

Step 5 On FusionInsight Manager, check whether the alarm whose ID is **12007**, **19000**, or **19013** and **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to [Step 8](#).
- If no, go to [Step 7](#).

Step 6 Run the **getusage /hbase/replication** command to check the capacity usage of the ZNode. Check whether the ratio of **Data size** to the ZNode capacity quota is close to the alarm threshold.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

Step 7 On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Check whether the alarm whose ID is **19006** and **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 9**.


Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 10 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.179 ALM-19018 HBase Compaction Queue Size Exceeds the Threshold

Alarm Description

The system checks the HBase compaction queue size every 30 seconds. This alarm is generated when the compaction queue size exceeds the alarm threshold (**100** by default) for three consecutive times. This alarm is cleared when the compaction queue size is less than the threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19018 | Minor | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |

Impact on the System

Write pressure on the HBase node continues going up, and the disk I/O and CPU may be overloaded. Read and write requests are responded slowly or even time out.

Possible Causes


- The number of HBase RegionServers is too small.
- There are excessive regions on a single RegionServer of HBase.
- The HBase RegionServer heap size is small.
- Resources are insufficient.
- Related parameters are not configured properly.

Handling Procedure

Check whether related parameters are properly configured.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. On the page that is displayed, check whether the alarm whose **Alarm ID** is **19008** or **19011** exists.
- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 3**.
 - If no, go to **Step 2**.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > HBase**. On the page that is displayed, click the **Configurations** tab then the **All Configurations** sub-tab, search for **hbase.hstore.compaction.min**, **hbase.hstore.compaction.max**, **hbase.regionserver.thread.compaction.small**, and **hbase.regionserver.thread.compaction.throttle**, and set them to larger values.
- Step 3** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to **Step 4**.

Collect the fault information.

- Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 5** Expand the **Service** drop-down list, and select **HBase** for the target cluster.
- Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 7** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.180 ALM-19019 Number of HBase HFiles to Be Synchronized Exceeds the Threshold**Alarm Description**

The system checks the number of HFiles to be synchronized by the RegionServer of each HBase service instance every 30 seconds. This indicator can be viewed on the RegionServer role monitoring page. This alarm is generated when the number of HFiles to be synchronized on a RegionServer exceeds the threshold (exceeding 128 for 20 consecutive times by default). To change the threshold, choose **O&M > Alarm > Threshold Configuration > Name of the desired cluster > HBase**. This alarm is cleared when the number of HFiles to be synchronized is less than or equal to the threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19019 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-----------|--|
| Source | Specifies the cluster for which the alarm was generated. |

| Parameter | Description |
|-------------------|--|
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

A large number of HFile files are stacked. Data is inconsistent between the active and standby nodes, and the latest data cannot be read from the standby cluster during an active/standby switchover or during HBase dual-read. If the fault persists, the storage space of the active cluster and ZooKeeper nodes will be used up. As a result, the active cluster service will be interrupted.

Possible Causes

- The network is abnormal.
- The RegionServer region distribution is unbalanced.
- The HBase service scale of the standby cluster is too small.

Handling Procedure

View alarm location information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19019**, and view the service instance and host name in **Location**.

Check the network connection between RegionServers on active and standby clusters.

Step 2 Run the **ping** command to check whether the network connection between the faulty RegionServer node and the host where RegionServer of the standby cluster resides is normal.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Contact the network administrator to restore the network.

Step 4 After the network recovers, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the RegionServer region distribution in the active cluster.

- Step 5** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**. Click **HMaster(Active)** to go to the web UI of the HBase instance and check whether regions are evenly distributed on the Region Server.

Region Servers

| ServerName | Start time | Last contact | Version | Requests Per Second | Num. Regions |
|-----------------------------------|--------------------------|--------------|------------------------------|---------------------|--------------|
| kwehpsra44947.21302.1620614446704 | 2021-05-10T02:40:46.704Z | 1 s | 2.2.3 hev-el-311001-SNAPSHOT | 13 | 10 |
| kwehpsra44949.21302.1620614361509 | 2021-05-10T02:39:21.509Z | 0 s | 2.2.3 hev-el-311001-SNAPSHOT | 0 | 12 |
| kwehpsra44949.21302.1620614361123 | 2021-05-10T02:39:21.123Z | 2 s | 2.2.3 hev-el-311001-SNAPSHOT | 0 | 13 |
| kwehpsr010223.21302.1621424421469 | 2021-05-19T11:40:21.469Z | 1 s | 2.2.3 hev-el-311001-SNAPSHOT | 0 | 8 |
| Total 4 | | | | 13 | 43 |

- Step 6** Log in to the faulty RegionServer node as user **omm**.

- Step 7** Run the following commands to go to the client installation directory and set the environment variable:

```
cd Client installation directory
```

```
source bigdata_env
```

If the cluster uses the security mode, perform security authentication. Run the **kinit hbase** command and enter the password as prompted (obtain the password from the MRS cluster administrator).

- Step 8** Run the following commands to check whether the load balancing function is enabled.

```
hbase shell
```

```
balancer_enabled
```

- If yes, go to [Step 10](#).
- If no, go to [Step 9](#).

- Step 9** Run the following commands in HBase Shell to enable the load balancing function and check whether the function is enabled.

```
balance_switch true
```

```
balancer_enabled
```

- Step 10** Run the **balancer** command to manually trigger the load balancing function.

NOTE

You are advised to enable and manually trigger the load balancing function during off-peak hours.

- Step 11** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

Check the HBase service scale of the standby cluster.

- Step 12** Expand the HBase cluster, add a node, and add a RegionServer instance on the node. Then, perform [Step 6](#) to [Step 10](#) to enable the load balancing function and manually trigger it.

- Step 13** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**. Click **HMaster(Active)** to go to the web UI of the HBase instance, refresh the page, and check whether regions are evenly distributed.

- If yes, go to [Step 14](#).
- If no, go to [Step 15](#).


Step 14 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 15](#).

Collect the fault information.

Step 15 On FusionInsight Manager of the standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 16 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 17 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.181 ALM-19020 Number of HBase WAL Files to Be Synchronized Exceeds the Threshold

Alarm Description

The system checks the number of WAL files to be synchronized by the RegionServer of each HBase service instance every 30 seconds. This indicator can be viewed on the RegionServer role monitoring page. This alarm is generated when the number of WAL files to be synchronized on a RegionServer exceeds the threshold (exceeding 128 for 20 consecutive times by default). To change the threshold, choose **O&M** > **Alarm** > **Threshold Configuration** > *Name of the desired cluster* > **HBase** . This alarm is cleared when the number of WAL files to be synchronized is less than or equal to the threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19020 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

A large number of WAL files are stacked. Data is inconsistent between the active and standby nodes, and the latest data cannot be read from the standby cluster during an active/standby switchover or during HBase dual-read. If the fault persists, the storage space of the active cluster and ZooKeeper nodes will be used up. As a result, the active cluster service will be interrupted.

Possible Causes

- The network is abnormal.
- The RegionServer region distribution is unbalanced.
- The HBase service scale of the standby cluster is too small.

Handling Procedure

View alarm location information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19020**, and view the service instance and host name in **Location**.

Check the network connection between RegionServers on active and standby clusters.

Step 2 Run the **ping** command to check whether the network connection between the faulty RegionServer node and the host where RegionServer of the standby cluster resides is normal.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Contact the network administrator to restore the network.

Step 4 After the network recovers, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check the RegionServer region distribution in the active cluster.

Step 5 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**. Click **HMaster(Active)** to go to the web UI of the HBase instance and check whether regions are evenly distributed on the Region Server.

Region Servers

| Base State | Memory | Requests | Storefiles | Compactions | Replications | | | | | |
|-----------------------------------|--------------------------|--------------|-----------------------------|---------------------|--------------|--|--|--|--|--|
| ServerName | Start time | Last contact | Version | Requests Per Second | Num. Regions | | | | | |
| kwehpsra44947.21302.1620614446704 | 2021-05-10T02:40:46.704Z | 1 s | 2.2.3-hw-el-311001-SNAPSHOT | 13 | 10 | | | | | |
| kwehpsra44949.21302.1620614361509 | 2021-05-10T02:39:21.509Z | 0 s | 2.2.3-hw-el-311001-SNAPSHOT | 0 | 12 | | | | | |
| kwehpsra44949.21302.1620614361123 | 2021-05-10T02:39:21.123Z | 2 s | 2.2.3-hw-el-311001-SNAPSHOT | 0 | 13 | | | | | |
| kwehpsr10223.21302.1621424421459 | 2021-05-10T11:40:21.459Z | 1 s | 2.2.3-hw-el-311001-SNAPSHOT | 0 | 8 | | | | | |
| Total 4 | | | | 13 | 43 | | | | | |

Step 6 Log in to the faulty RegionServer node as user **omm**.

Step 7 Run the following commands to go to the client installation directory and set the environment variable:

```
cd Client installation directory
```

```
source bigdata_env
```

If the cluster uses the security mode, perform security authentication. Run the **kinit hbase** command and enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 8 Run the following commands to check whether the load balancing function is enabled.

```
hbase shell
```

```
balancer_enabled
```

- If yes, go to [Step 10](#).
- If no, go to [Step 9](#).

Step 9 Run the following commands in HBase Shell to enable the load balancing function and check whether the function is enabled.

```
balance_switch true
```

```
balancer_enabled
```

Step 10 Run the **balancer** command to manually trigger the load balancing function.


NOTE

You are advised to enable and manually trigger the load balancing function during off-peak hours.

Step 11 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

Check the HBase service scale of the standby cluster.

- Step 12** Expand the HBase cluster, add a node, and add a RegionServer instance on the node. Then, perform [Step 6](#) to [Step 10](#) to enable the load balancing function and manually trigger it.
- Step 13** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**. Click **HMaster(Active)** to go to the web UI of the HBase instance, refresh the page, and check whether regions are evenly distributed.
- If yes, go to [Step 14](#).
 - If no, go to [Step 15](#).
- Step 14** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 15](#).
- Collect the fault information.**
- Step 15** On FusionInsight Manager of the standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.
- Step 16** Expand the **Service** drop-down list, and select **HBase** for the target cluster.
- Step 17** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 18** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.182 ALM-19021 Handler Usage of RegionServer Exceeds the Threshold

Alarm Description

The system checks the RegionServer handler usage of each HBase service instance every 30 seconds. This alarm is generated when the handler usage of a RegionServer exceeds the threshold (90% for five consecutive times by default). This alarm is cleared if the handler usage is lower than or equal to the threshold.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19021 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

RegionServer may not be able to provide services externally. As a result, the concurrent read/write processing become slow, or requests fail.

Possible Causes

- The value of a handler is too small.
- Hotspotting occurs.

Handling Procedure

View alarm location information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**, locate the row that contains the alarm whose **Alarm ID** is **19021**, and view the service instance and host name in **Location**.

Check the handler configuration.

Step 2 Choose **Cluster > Services > HBase** and click the **Configurations** tab. In the upper right corner of the page, search for **hbase.regionserver.handler.count** and check whether its value is too small. The default value is **200**.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Change the value of this parameter to a larger value and save the configuration. Choose **Cluster > Services > HBase**, click the **Instance** tab, select the affected RegionServer instances, and choose **More > Instance Rolling Restart**. In the displayed dialog box, enter the username and password. In the **Instance Rolling Restart** dialog box, click **OK** and wait until the rolling restart is complete.

Step 4 After the configuration takes effect, check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to [Step 5](#).

Check whether cluster hotspotting occurs.

Step 5 On FusionInsight Manager, choose **Cluster > Services > HBase** and click **HMaster(Active)** following **HMaster WebUI** to go to the web UI of the HBase instance. In the **Region Servers** area of the **Home** page, click **Requests** and check whether the requests in the **Filtered Read Request Count** and **Write Request Count** columns are evenly distributed.

Region Servers

| ServerName | Request Per Second | Read Request Count | Filtered Read Request Count | Write Request Count |
|------------|--------------------|--------------------|-----------------------------|---------------------|
| | 0 | 4591 | 0 | 1460 |
| | 0 | 708601 | 1957 | 1375 |
| | 0 | 3472032 | 685564 | 1183 |

- If yes, go to [Step 13](#).
- If no, go to [Step 6](#).

Step 6 Check whether regions are evenly distributed.

On FusionInsight Manager, choose **Cluster > Services > HBase** and click **HMaster(Active)** following **HMaster WebUI** to go to the web UI of the HBase instance. In the **Region Servers** area of the **Home** page, click **Base Stats** and check whether the regions in the **Num.Regions** column are evenly distributed.

Region Servers

| ServerName | Start time | Last contact | Version | Requests Per Second | Num. Regions |
|------------|--------------------------|--------------|---------|---------------------|--------------|
| | 2021-05-10T02:40:46.704Z | 1 s | | 13 | 10 |
| | 2021-05-10T02:39:21.509Z | 0 s | | 0 | 12 |
| | 2021-05-10T02:39:21.123Z | 2 s | | 0 | 13 |
| | 2021-05-10T11:40:21.456Z | 1 s | | 0 | 6 |
| Total:4 | | | | 13 | 43 |

- If yes, go to [Step 13](#).
- If no, go to [Step 7](#).

Step 7 Log in to the faulty RegionServer node as user **omm**.

Step 8 Run the following commands to go to the client installation directory and set the environment variable:

```
cd Client installation directory
```

```
source bigdata_env
```

If the cluster uses the security mode, run the following command to perform security authentication:

```
kinit hbase
```

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 9 Run the following commands to check whether the load balancing function is enabled. If the command output is **true**, the load balancing function is enabled.

```
hbase shell
```

```
balancer_enabled
```

```
hbase:004:0> balancer_enabled
true
```

```
Took 0.0165 seconds  
=> true
```

- If yes, go to [Step 13](#).
- If no, go to [Step 10](#).

Step 10 Run the following commands in HBase Shell to enable the load balancing function and check whether the function is enabled.

```
balance_switch true
```

```
balancer_enabled
```

 **NOTE**

You are advised to enable and manually trigger the load balancing function during off-peak hours.

Step 11 Run the following command to manually trigger the load balancing function:

```
balancer
```


Step 12 After the load balancing is complete, log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 13](#).

Collect the fault information.

Step 13 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 14 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 15 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 16 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.183 ALM-19022 HBase Hotspot Detection Is Unavailable

Alarm Description

When the MetricController instance is installed for HBase, the alarm module checks the health status of the active HBase MetricController instance every 120

seconds. This alarm is generated when the active HBase MetricController instance does not exist or is unavailable and the hotspot detection function is unavailable.

This alarm is cleared when the active HBase MetricController instance recovers.

 **NOTE**

This alarm applies only to MRS 3.3.0 or later.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19022 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

The HBase hotspot detection function is unavailable. Services are not affected. However, if request/data skew occurs, the system cannot report alarms and automatically recovers from hotspotting. Service requests may cause node overload, slow response, and request timeout.

Possible Causes

- The ZooKeeper service is abnormal.
- The HBase service is abnormal.
- In the current HBase service, the MetricController instance on the same node as the active HMaster instance is not started.
- The network is abnormal.

Handling Procedure

Check the ZooKeeper service status.

- Step 1** In the service list on FusionInsight Manager, check whether **Running Status** of ZooKeeper is **Normal**.
- If yes, go to [Step 5](#).
 - If no, go to [Step 2](#).
- Step 2** In the alarm list, check whether **ALM-13000 ZooKeeper Service Unavailable** exists.
- If yes, go to [Step 3](#).
 - If no, go to [Step 5](#).
- Step 3** Rectify the fault by performing the operations provided for **ALM-13000 ZooKeeper Service Unavailable**.
- Step 4** Wait for several minutes and check whether the alarm **HBase Hotspot Detection Is Unavailable** is cleared.
- If yes, no further action is required.
 - If no, go to [Step 5](#).
- Check the HBase service status.**
- Step 5** In the service list on FusionInsight Manager, check whether **Running Status** of HBase is **Normal**.
- If yes, go to [Step 9](#).
 - If no, go to [Step 6](#).
- Step 6** In the alarm list, check whether the alarm **ALM-19000 HBase Service Unavailable** exists.
- If yes, go to [Step 7](#).
 - If no, go to [Step 9](#).
- Step 7** Rectify the fault by following the steps provided for **ALM-19000 HBase Service Unavailable**.
- Step 8** Wait for several minutes and check whether the alarm **HBase Hotspot Detection Is Unavailable** is cleared.
- If yes, no further action is required.
 - If no, go to [Step 9](#).
- Check whether the MetricController instance deployed on the same node as the active HMaster instance is started.**
- Step 9** On FusionInsight Manager, choose **Cluster > Service > HBase**, and click **Instances** to check whether the **MetricController(Active)** instance exists.
- If yes, go to [Step 12](#).
 - If no, go to [Step 10](#).
- Step 10** Select the MetricController instance whose management IP address is the same as that of the active HMaster instance, and click **Start Instance**.
- Step 11** After the MetricController instance is restarted, check whether the alarm **HBase Hotspot Detection Is Unavailable** is cleared.
- If yes, no further action is required.

- If no, go to [Step 12](#).

Check the network connectivity between the started MetricController instances and the active HMaster node.

Step 12 Log in to the node where the active HMaster instance is deployed and run **ping** *IP address of the node where the standby MetricController instance is deployed* to check whether the network connection between the started MetricController instances and the host where the active HMaster instance is deployed is normal.

- If yes, go to [Step 15](#).
- If no, go to [Step 13](#).

Step 13 Contact the network administrator to restore the network.

Step 14 After the network recovers, check whether the alarm **HBase Hotspot Detection Is Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 15](#).

Collect fault information.

Step 15 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 16 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 17 In the **Host** area, select the host where the HMaster instance is deployed.

Step 18 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.184 ALM-19023 Region Traffic Restriction for HBase

Alarm Description

When the MetricController instance is installed for the HBase service, self-healing from hotspotting is automatically enabled. The alarm module checks whether there are regions whose request traffic is restricted due to hotspot issues in HBase every 120 seconds. This alarm is generated when the region where hotspot traffic is restricted is detected in HBase.

This alarm is cleared when the region is no longer a hotspot.

This alarm applies only to MRS 3.3.0 or later.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19023 | Critical | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

If the traffic of a hotspot region is restricted, the number of handlers for processing the requests in the region is limited. As a result, services requesting the region may slow down or retry upon failure.

Possible Causes

Too many requests are directed to a single region when the HBase service is accessed.

Handling Procedure

Check whether there are too many requests in a single region of HBase.

- Step 1** Log in to FusionInsight Manager, and Choose **O&M > Alarm > Alarms**.
- Step 2** In **Additional Information** of **Region Traffic Restriction for HBase**, view the reported table name and region information.
- Step 3** On FusionInsight Manager, choose **Cluster > Service > HBase** and click the hyperlink on the right of HMaster web UI.
- Step 4** Click **Table Details** and adjust service configurations in the region where the table in **Step 2** is deployed.
- Step 5** Wait a moment and then check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 7 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 8 In the **Host** area, select the host where the HMaster instance is deployed.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm will be automatically cleared.

Related Information

None.

7.12.185 ALM-19024 RPC Requests P99 Latency on RegionServer Exceeds the Threshold

Alarm Description

The system checks P99 latency for RPC requests on each RegionServer instance of the HBase service every 30 seconds. This alarm is generated when P99 latency for RPC requests on a RegionServer exceeds the threshold for 10 consecutive times.

This alarm is cleared when P99 latency for RPC requests on a RegionServer instance is less than or equal to the threshold.

This alarm applies only to MRS 3.3.0 or later.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|---|--------------|
| 19024 | <ul style="list-style-type: none">• Critical: The default threshold is 10 seconds.• Major: The default threshold is 5 seconds. | Yes |

Alarm Parameters

| Parameter | Description |
|-------------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm. |

Impact on the System

If RPC requests P99 latency exceeds the threshold, the RegionServer cannot deliver normal service performance externally. For latency-sensitive services, a large number of service read and write requests may time out.

Possible Causes

- RegionServer GC duration is too long.
- The HDFS RPC response is too slow.
- RegionServer request concurrency is too high.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19024**, and view the service instance and host name in **Location**.

Check the GC duration of RegionServer.

Step 2 In the alarm list on FusionInsight Manager, check whether the "HBase GC Duration Exceeds the Threshold" alarm is generated for the service instance in **Step 1**.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by following the handling procedure of "ALM-19007 HBase GC Duration Exceeds the Threshold".

Step 4 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check HDFS RPC response time.

Step 5 In the alarm list on FusionInsight Manager, check whether alarm "Average NameNode RPC Processing Time Exceeds the Threshold" is generated for the HDFS service on which the HBase service depends.

- If yes, go to [Step 6](#).
- If no, go to [Step 8](#).

Step 6 Rectify the fault by following the handling procedure of "ALM-14021 Average NameNode RPC Processing Time Exceeds the Threshold".

Step 7 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

Check the number of concurrent processes on a RegionServer.

Step 8 In the alarm list on FusionInsight Manager, check whether the "Handler Usage of RegionServer Exceeds the Threshold" alarm is generated for the service instance in [Step 1](#).

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

Step 9 Rectify the fault by following the handling procedure of "ALM-19021 Handler Usage of RegionServer Exceeds the Threshold".

Step 10 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

Collect fault information.

Step 11 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 12 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 13 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.186 ALM-19025 Damaged StoreFile in HBase

Alarm Description

The system checks the **hdfs://hacluster/hbase/autocorrupt** and **hdfs://hacluster/hbase/MasterData/autocorrupt** directories on HDFS of each HBase service every 120 seconds. This alarm is generated when there are files in the directories.

This alarm is cleared when the **hdfs://hacluster/hbase/autocorrupt** and **hdfs://hacluster/hbase/MasterData/autocorrupt** directories do not exist or are empty.

This alarm applies only to MRS 3.3.0 or later.

NOTE

hdfs://hacluster indicates the name of the file system used by HBase, and **/hbase** indicates the root directory of HBase in the file system. You can log in to FusionInsight Manager, choose **Cluster > Services > HBase** and click **Configuration**. Search for **fs.defaultFS** and **hbase.data.rootdir**.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 19025 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|---|
| Source | Specifies the cluster for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System

Data in the file may be lost, and data queried by the service may be inconsistent.

Possible Causes

The StoreFile files are damaged.

Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19025**, and view the service in **Location**.
- Step 2** Log in to the node where the HDFS and HBase clients are installed as the client installation user and run the following commands:
- ```
cd Client installation directory
source bigdata_env
kinit Component service user (If Kerberos authentication is disabled for the cluster (the cluster is in normal mode), skip this step.)
```
- Step 3** Check the damaged StoreFile file.
- Run the following command to check whether the **/hbase/autocorrupt** directory of HDFS is empty. If it is not, go to **Step 4**.  
**hdfs dfs -ls -R hdfs://hacluster/hbase/autocorrupt**
  - Run the following command to check whether the **/hbase/MasterData/autocorrupt** directory of HDFS is empty. If it is not, go to **Step 9**.  
**hdfs dfs -ls -R hdfs://hacluster/hbase/MasterData/autocorrupt**
- Step 4** Run the following command to restore the StoreFile files in the **hdfs://hacluster/hbase/autocorrupt** directory:
- ```
hdfs debug recoverLease -path hdfs://hacluster/hbase/autocorrupt/Name space/Table/Region/Column family/StoreFile files
```
- Step 5** Check whether the damaged StoreFile files are restored. If the following information is displayed, the restoration is successful:
- ```
recoverLease SUCCEEDED on hdfs://hacluster/hbase/autocorrupt/
default/h1/865665fe32db62dadada68b644359809/cf1/95f210f931ad44c99e4028470be7d292
```
- If yes, go to **Step 6**.  
If no, go to **Step 9**.
- Step 6** Run the following command to move the files back to the **hdfs://hacluster/hbase/data** directory:
- ```
hdfs dfs -mv hdfs://hacluster/hbase/autocorrupt/Name space/Table/Region/  
Column family/StoreFile fileshdfs://hacluster/hbase/data/Name space/Table/  
Region/Column family/StoreFile files
```
- Step 7** Run the following command on HBase Shell to bring the region online again:
- ```
hbase shell
unassign'Region'
assign'Region'
```
- Step 8** Wait several minutes and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 9**.
- Collect fault information.**



- Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
  - Step 10** Expand the **Service** drop-down list, and select **HBase** for the target cluster.
  - Step 11** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
  - Step 12** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

# 7.12.187 ALM-19026 Damaged WAL Files in HBase

## Alarm Description

The system checks the **hdfs://hacluster/hbase/corrupt** directory on the HDFS of each HBase service every 120 seconds. This alarm is generated when there are WAL files in the **/hbase/corrupt** directory.

This alarm is cleared when the **/hbase/corrupt** directory does not exist or does not contain WAL files.

This alarm applies only to MRS 3.3.0 or later.

### NOTE

**hdfs://hacluster** indicates the name of the file system used by HBase, and **/hbase** indicates the root directory of HBase in the file system. You can log in to FusionInsight Manager, choose **Cluster > Services > HBase** and click **Configuration**. Search for **fs.defaultFS** and **hbase.data.rootdir**.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19026	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

If the data in the damaged file is not flushed to disks, the data will be lost. As a result, some data queried by the service is inconsistent.

## Possible Causes

The WAL files are damaged.

## Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19026**, and view the service in **Location**.
  - Step 2** Log in to the node where the HDFS clients are installed as the client installation user and run the following commands:  

```
cd Client installation directory
```

```
source bigdata_env
```

```
kinit Component service user
```

(If Kerberos authentication is disabled for the cluster (the cluster is in normal mode), skip this step.)
  - Step 3** Run the following command to check the damaged WAL files and go to **Step 4**:  

```
hdfs dfs -ls hdfs://hacluster/hbase/corrupt/*%2C*
```

**Collect fault information.**
  - Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
  - Step 5** Expand the **Service** drop-down list, and select **HBase** for the target cluster.
  - Step 6** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
  - Step 7** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.188 ALM-20002 Hue Service Unavailable

### Description

This alarm is generated when the Hue service is unavailable. The system checks the Hue service status every 60 seconds.

This alarm is cleared when the Hue service is normal.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
20002	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

Users cannot perform interactive analysis and data processing with MRS on the Hue UI.

### Possible Causes

- The internal KrbServer service on which the Hue service depends is abnormal.
- The internal DBService service on which the Hue service depends is abnormal.
- The network connection to the DBService is abnormal.

## Procedure

### Check whether the KrbServer is abnormal.

- Step 1** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services**. In the service list, check whether the **KrbServer** running status is **Normal**.
- If yes, go to [Step 4](#).
  - If no, go to [Step 2](#).

**Step 2** Restart the KrbServer service.

- Step 3** Wait several minutes, and check whether **Hue Service Unavailable** is cleared.
- If yes, no further action is required.
  - If no, go to [Step 4](#).

### Check whether the DBService is abnormal.

- Step 4** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services**.
- Step 5** In the service list, check whether the **DBService** running status is **Normal**.
- If yes, go to [Step 8](#).
  - If no, go to [Step 6](#).
- Step 6** Restart the DBService.

---

#### NOTICE

When the service is rebooted, it becomes unavailable and can disrupt business operations.

---

- Step 7** Wait several minutes, and check whether **Hue Service Unavailable** is cleared.
- If yes, no further action is required.
  - If no, go to [Step 8](#).

### Check whether the network connection to the DBService is normal.

- Step 8** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hue** > **Instance**, record the IP address of the active Hue.
- Step 9** Log in to the active Hue.
- Step 10** Run the **ping** command to check whether communication between the host that runs the active Hue and the hosts that run the DBService is normal. (Obtain the IP addresses of the hosts that run the DBService in the same way as that for obtaining the IP address of the active Hue.)
- If yes, go to [Step 13](#).
  - If no, go to [Step 11](#).
- Step 11** Contact the administrator to restore the network.
- Step 12** Wait several minutes, and check whether **Hue Service Unavailable** is cleared.


- If yes, no further action is required.
- If no, go to [Step 13](#).

**Collect fault information.**

**Step 13** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 14** Select the following nodes in the required cluster from the **Service** drop-down list:

- Hue
- Controller

**Step 15** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 16** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.189 ALM-23001 Loader Service Unavailable

### Description

The system checks the Loader service availability every 60 seconds. This alarm is generated when the system detects that the Loader service is unavailable. This alarm is cleared when the Loader service is available.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23001	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

When the Loader service is unavailable, the data loading, import, and conversion functions are unavailable.

## Possible Causes

- The internal service on which the Loader service depends is abnormal.
  - The ZooKeeper service is abnormal.
  - The HDFS service is abnormal.
  - The DBService service is abnormal.
  - The Yarn service is abnormal.
  - The Mapreduce service is abnormal.
- Environment fault: The network is abnormal, which the Loader service cannot communicate with the depended internal services and cannot provide services.
- Software fault: The Loader service cannot run properly.

## Procedure

### Check the ZooKeeper service status.

- Step 1** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** to check whether the ZooKeeper running status is **Normal**.
- If yes, go to [Step 3](#).
  - If no, go to [Step 2](#).
- Step 2** Choose **More** > **Restart Service** to restart the ZooKeeper service. In the alarm list, check whether **LoaderService Unavailable** is cleared.
- If yes, no further action is required.
  - If no, go to [Step 3](#).
- Step 3** On the FusionInsight Manager, check whether the alarm list contains **Process Fault**.
- If yes, go to [Step 4](#).
  - If no, go to [Step 7](#).
- Step 4** In the **Location** area of **Process Fault**, check whether **ServiceName** is **ZooKeeper**.
- If yes, go to [Step 5](#).

- If no, go to [Step 7](#).

**Step 5** Rectify the fault by following the steps provided in **ALM-12007 Process Fault**.

**Step 6** In the alarm list, check whether **Loader Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Check the HDFS service status.**

**Step 7** On the FusionInsight Manager, check whether the alarm list contains **HDFS Service Unavailable**.

- If yes, go to [Step 8](#).
- If no, go to [Step 10](#).

**Step 8** Rectify the fault by following the steps provided in **ALM-14000 HDFS Service Unavailable**.

**Step 9** In the alarm list, check whether **Loader Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

**Check the DBService status.**

**Step 10** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService** to check whether the DBService running status is **Normal**.

- If yes, go to [Step 12](#).
- If no, go to [Step 11](#).

**Step 11** Choose **More** > **Restart Service** to restart the DBService service. In the alarm list, check whether **LoaderService Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

**Check the Mapreduce status.**

**Step 12** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Mapreduce** to check whether the Mapreduce running status is **Normal**.

- If yes, go to [Step 16](#).
- If no, go to [Step 13](#).

**Step 13** Choose **More** > **Restart Service** to restart the Mapreduce service. In the alarm list, check whether **LoaderService Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 16](#).

**Check the Yarn status.**

**Step 14** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** to check whether the Yarn running status is **Normal**.

- If yes, go to [Step 16](#).

- If no, go to [Step 15](#).

**Step 15** Choose **More > Restart Service** to restart the Yarn service. In the alarm list, check whether **LoaderService Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 16](#).

**Step 16** On the FusionInsight Manager, check whether the alarm list contains **Yarn Service Unavailable**.

- If yes, go to [Step 17](#).
- If no, go to [Step 19](#).

**Step 17** Rectify the fault by following the steps provided in **ALM-18000 Yarn Service Unavailable**.

**Step 18** In the alarm list, check whether **Loader Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 19](#).

**Check the network connection between Loader and dependent components.**

**Step 19** On the FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Loader**.

**Step 20** Click **Instance** and the LoaderServer instance list is displayed.

**Step 21** Record the **Management IP Address** in the row of **LoaderServer(Active)**.

**Step 22** Log in to the host where the active LoaderServer runs as **omm** user using the IP address obtained in [Step 21](#).

**Step 23** Run the **ping** command to check whether communication between the host that runs the active LoaderServer and the hosts that run the dependent components. (The dependent components include ZooKeeper, DBService, HDFS, Mapreduce and Yarn. Obtain the IP addresses of the hosts that run these services in the same way as that for obtaining the IP address of the active LoaderServer.)

- If yes, go to [Step 26](#).
- If no, go to [Step 24](#).

**Step 24** Contact the administrator to restore the network.

**Step 25** In the alarm list, check whether **Loader Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 26](#).

**Collect fault information.**


**Step 26** On the FusionInsight Manager, choose **O&M > Log > Download**.

**Step 27** Select the following nodes in the required cluster from the **Service** drop-down list:

- ZooKeeper
- HDFS
- DBService



- Yarn
- Mapreduce
- Loader

**Step 28** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 29** On the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.

**Step 30** Choose **More** > **Restart Service**, and click **OK**.

**Step 31** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 32](#).

**Step 32** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.190 ALM-23003 Loader Task Execution Failure

### Description

This alarm is generated immediately when the system detects that the Loader job fails. This alarm is cleared when the failed job is manually handled by a user. This alarm must be manually cleared.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23003	Minor	No

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
JobID	Specifies the ID of failed Loader job.
JobName	Specifies the failed Loader job.
UserName	Specifies the name of the user who submits the Loader job.
Details	Supplementary information for which the alarm is generated.

## Impact on the System

This is a job-level alarm for Loader. The job execution fails, and you need to view specific logs to locate the failure cause. No execution result is returned. After the fault is rectified, you need to execute the task again. No impact on the Loader service.

## Possible Causes

- Task parameters are incorrectly configured.
- Exceptions occur when Yarn is executing a job.

## Procedure

### Check whether task parameters are incorrectly configured.

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms** and click the alarm drop-down list from the alarm list, obtain the **Alarm Cause**.
- Step 2** If the alarm cause is "Failure to submit job", view error details in **Additional Information**, and go to the Loader WebUI to view the execution history of the job.

#### NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

- Step 3** Submit the task again.
- Step 4** Check whether the task executed successfully.
- If yes, go to [Step 9](#).

- If no, go to [Step 5](#).

**Check whether exceptions occur when Yarn is executing a job.**

**Step 5** On FusionInsight Manager, click the alarm drop-down list from the alarm list, obtain the **Alarm Cause**.

**Step 6** Check whether the Yarn activity is executed properly in the **Alarm Cause**. If the alarm cause is "Yarn execution failed", the Yarn activity is abnormal.

- If yes, go to [Step 7](#).
- If no, go to [Step 10](#).

**Step 7** Submit the task again.

**Step 8** Please check whether the task executed successfully.

- If yes, go to [Step 9](#).
- If no, go to [Step 10](#).


**Step 9** In the alarm list, click **Clear** from **Operation** to manually clear the alarm. No further action is required.

**Collect fault information.**

**Step 10** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 11** Select the following nodes in the required cluster from the **Service** drop-down list:

- DBService
- HDFS
- Loader
- Mapreduce
- Yarn
- ZooKeeper

**Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system does not automatically clear this alarm, and you need to manually clear the alarm.

## Related Information

None

## 7.12.191 ALM-23004 Loader Heap Memory Usage Exceeds the Threshold

### Description

The system checks the heap memory usage of the Loader service every 60 seconds. The alarm is generated when the heap memory usage of a Loader instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. The alarm is cleared when the heap memory usage is less than the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23004	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

### Impact on the System

Full GC occurs frequently in Loader. The performance deteriorates and page responses are slow. If the memory overflows, Loader may fail to provide services for external systems. The Loader page cannot be accessed, interfaces cannot be called, and active/standby switchover is frequently performed due to exceptions.

### Possible Causes

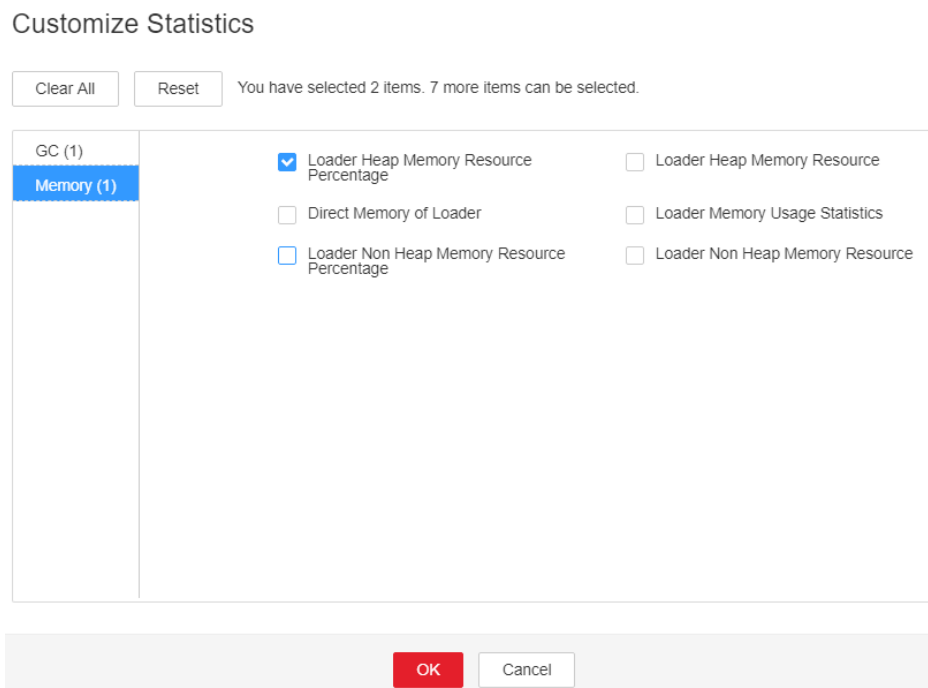
The heap memory of the Loader instance is overused or the heap memory is inappropriately allocated.

## Procedure

### Check heap memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Loader Heap Memory Usage Exceeds the Threshold > Location**. Check the host name of the instance involved in this alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Loader > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize > Memory > Loader Heap Memory Resource Percentage**. Click **OK**.

**Figure 7-116** Loader Heap Memory Resource Percentage



- Step 3** Check whether the used heap memory of Loader reaches the threshold (the default value is 95% of the maximum heap memory) specified for Loader.
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.
- Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Loader > Configurations**. Click **All Configurations**. Increase the value of **-Xmx** in **GC\_OPTS** as required, and click **Save**. Click **OK**.

 NOTE

- If this alarm is generated, the heap memory configured for the current Loader instance is insufficient for data transmission. You are advised to open the instance monitoring page, display the Loader heap memory resource status monitoring chart, and observe the change trend of the heap memory used by Loader in the monitoring chart. Then change the value of **-Xmx** to twice the current heap memory usage or to another value to meet site requirements.
- When setting the heap memory, you can set **-Xms** and **-Xmx** to similar values to avoid performance deterioration caused by heap size adjustment after each GC.
- Ensure that the sum of **-Xmx** and **XX:MaxPermSize** is not greater than the physical memory of the node server.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).


 NOTE

During service or instance restart, Loader cannot provide services to external systems. New jobs cannot be submitted, but jobs being executed are not affected.

**Collect fault information.**

**Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 7** Select **Loader** in the required cluster from the **Service** drop-down list.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected fault logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.192 ALM-23005 Loader Non-Heap Memory Usage Exceeds the Threshold

### Description

The system checks the non-heap memory usage of the Loader service every 30 seconds. The alarm is generated when the non-heap memory usage of a Loader instance exceeds the threshold (80% of the maximum memory) for 5 consecutive times. The alarm is cleared when the non-heap memory usage is less than the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23005	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

The Loader page may fail to be accessed and cannot provide services for external systems.

## Possible Causes

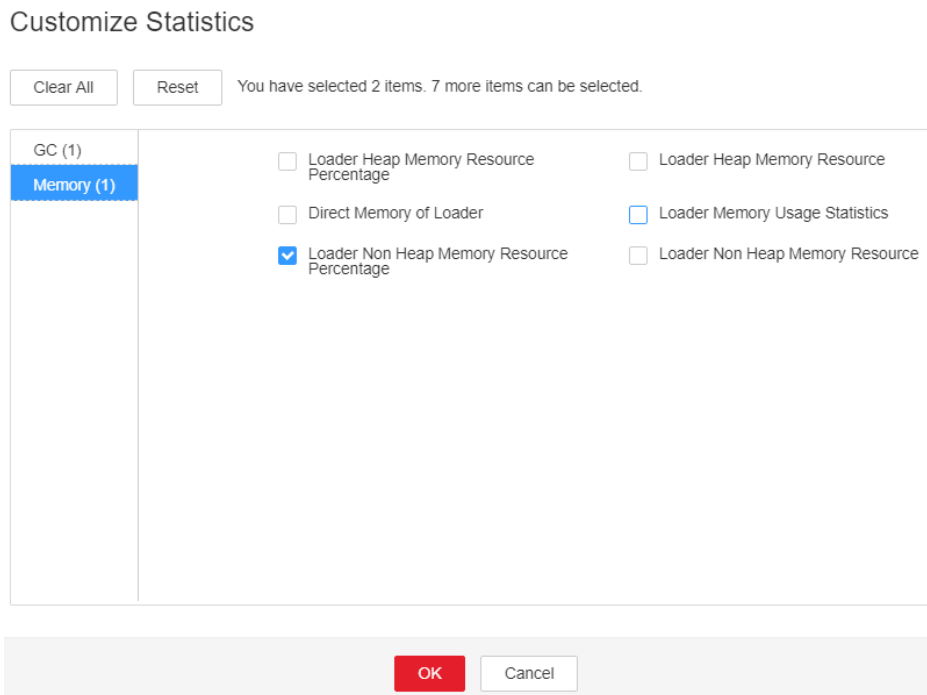
The non-heap memory of the Loader instance is overused or the non-heap memory is inappropriately allocated.

## Procedure

**Check non-heap memory usage.**

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Loader Non-Heap Memory Usage Exceeds the Threshold > Location**. Check the host name of the instance involved in this alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Loader > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize > Memory > Loader Non Heap Memory Resource Percentage**. Click **OK**.

**Figure 7-117** Loader Non Heap Memory Resource Percentage



**Step 3** Check whether the used non-heap memory of Loader reaches the threshold (the default value is 80% of the maximum non-heap memory) specified for Loader.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Loader > Configurations**. Click **All Configurations** Search **LOADER\_GC\_OPTS** in the search box. If the **-XX:MaxPermSize** parameter is not configured, set the initial value to **-XX:MaxPermSize=256M** for the first time. (If the alarm persists after the first adjustment, perform the second adjustment by referring to the following note.) And click **Save**. Click **OK**.

**NOTE**

If this alarm is generated, the non-heap memory configured for the current Loader instance is insufficient for the service scenario. You are advised to open the instance monitoring page, open the Loader non-heap memory resource status monitoring chart, and observe the change trend of the non-heap memory used by Loader in the monitoring chart. Then change the value of **-XX:MaxPermSize** to twice the current non-heap memory usage or to another value to meet site requirements.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to **Step 6**.

**NOTE**

During service or instance restart, Loader cannot provide services to external systems. New jobs cannot be submitted, but jobs being executed are not affected.

**Collect fault information.**



- Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 7** Select **Loader** in the required cluster from the **Service** drop-down list.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected fault logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.193 ALM-23006 Loader Direct Memory Usage Exceeds the Threshold

## Description

The system checks the direct memory usage of the Loader service every 30 seconds. The alarm is generated when the direct memory usage of a Loader instance exceeds the threshold (80% of the maximum memory) for 5 consecutive times. The alarm is cleared when the direct memory usage of Loader is less than or equal to the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23006	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

Loader may fail to provide services for external systems. I/O or socket exceptions occur, and active/standby switchovers occur frequently.

## Possible Causes

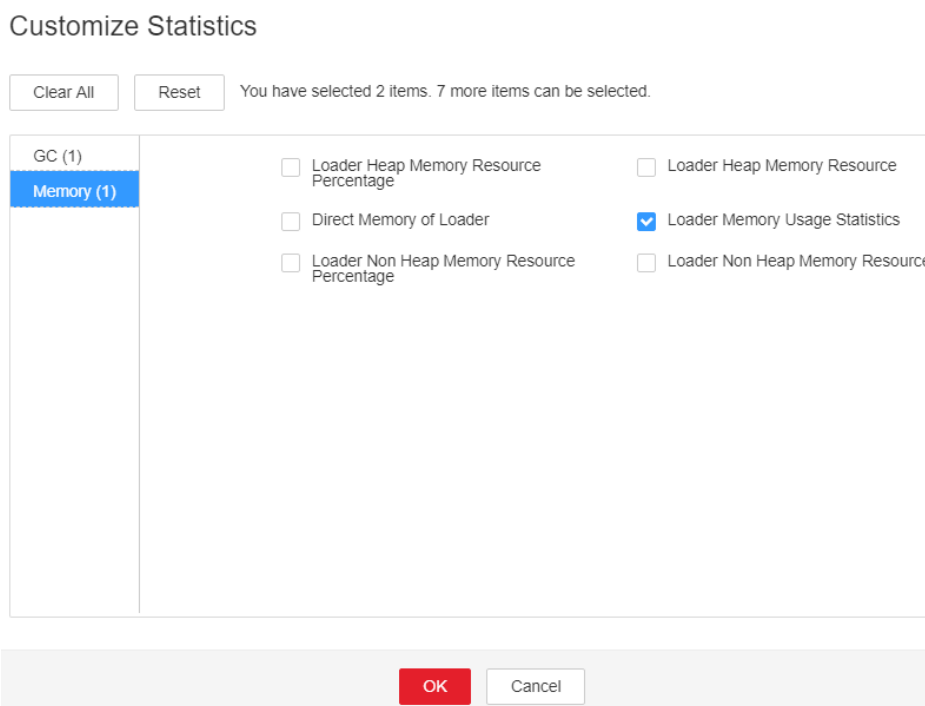
The direct memory of the Loader instance is overused or the direct memory is inappropriately allocated.

## Procedure

**Check direct memory usage.**

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Loader Direct Memory Usage Exceeds the Threshold > Location**. Check the host name of the instance involved in this alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Loader > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize > Memory > Loader Memory Usage Statistics**. Click **OK**.

**Figure 7-118** Loader Memory Usage Statistics



- Step 3** Check whether the used direct memory of Loader reaches the threshold (the default value is 80% of the maximum direct memory) specified for Loader.
- If yes, go to [Step 4](#).
  - If no, go to [Step 6](#).
- Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Configurations**. Click **All Configurations**. Search **LOADER\_GC\_OPTS** in the search box. Increase the value of **-XX:MaxDirectMemorySize** as required, and click **Save**. Click **OK**.

 **NOTE**


If this alarm is generated, the direct memory configured for the current Loader instance is insufficient for the service scenario. You are advised to open the instance monitoring page, display the Loader direct memory resource status monitoring chart, and observe the change trend of the direct memory used by Loader in the monitoring chart. Then change the value of **-XX:MaxDirectMemorySize** to twice the current direct memory usage or to another value to meet site requirements.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).

 **NOTE**

During service or instance restart, Loader cannot provide services to external systems. New jobs cannot be submitted, but jobs being executed are not affected.

**Collect fault information.**

- Step 6** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.
- Step 7** Select **Loader** in the required cluster from the **Service** drop-down list.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected fault logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.194 ALM-23007 Garbage Collection (GC) Time of the Loader Process Exceeds the Threshold

## Description

The system checks GC time of the Loader process every 60 seconds. The alarm is generated when GC time of the Loader process exceeds the threshold (default

value: **12 seconds**) for 5 consecutive times. The alarm is cleared when GC time is less than the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23007	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

Full GC occurs frequently, and the Loader service responds slowly. The Loader service may even break down and cannot provide services properly.

## Possible Causes

The heap memory of the Loader instance is overused or the heap memory is inappropriately allocated.

## Procedure

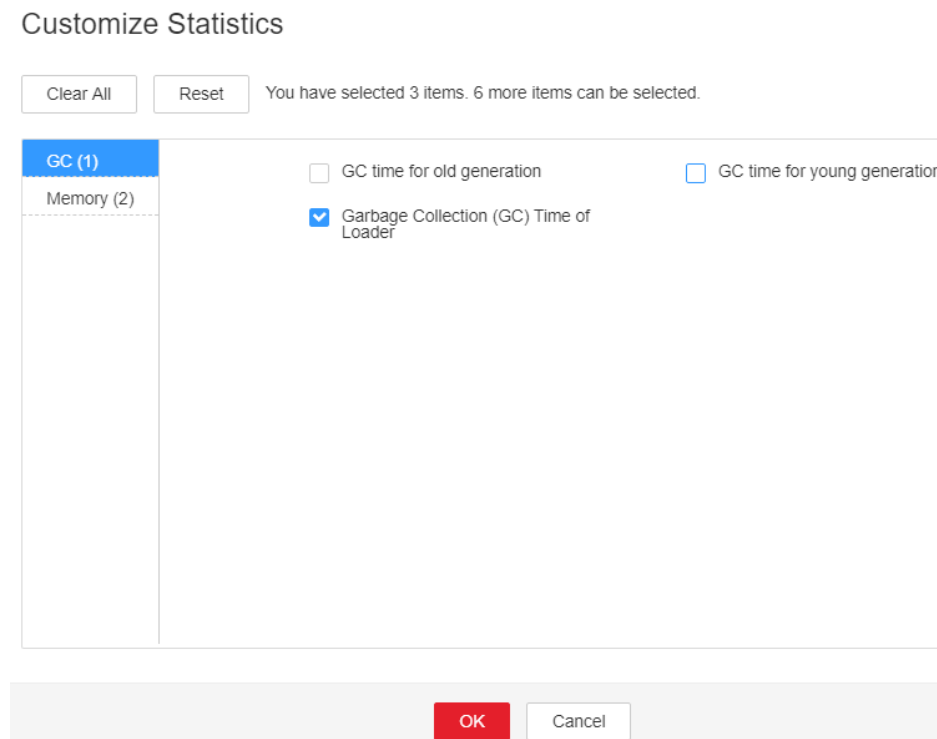
**Check GC time.**

**Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Garbage Collection (GC) Time of the Loader Process Exceeds the Threshold > Location**. Check the host name of the instance involved in this alarm.

**Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Loader > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the

chart area and choose **Customize > GC > Garbage Collection (GC) Time of Loader**. Click **OK**.

**Figure 7-119** Garbage Collection (GC) Time of Loader



**Step 3** Check whether GC time of the Loader process every second exceeds the threshold (default value: **12 seconds**).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Loader > Configurations**. Click **All Configurations**. Search **LOADER\_GC\_OPTS** in the search box. Increase the value of **-Xmx** as required, and click **Save**. Click **OK**.

**NOTE**

If this alarm is generated, the heap memory configured for the current Loader instance cannot meet the heap memory required for data transmission. You are advised to handle the problem by referring to **Step 4** in section **ALM-23004 Loader Heap Memory Usage Exceeds the Threshold**.


**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

**NOTE**

During service or instance restart, Loader cannot provide services to external systems. New jobs cannot be submitted, but jobs being executed are not affected.

**Collect fault information.**

- Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 7** Select **Loader** in the required cluster from the **Service** drop-down list.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected fault logs.
- End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.195 ALM-24000 Flume Service Unavailable

### Alarm Description

The alarm module checks the Flume service status every 180 seconds. This alarm is generated if the Flume service is abnormal.

This alarm is automatically cleared after the Flume service recovers.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24000	Critical	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

Flume cannot work and data transmission is interrupted.

## Possible Causes

All Flume instances are faulty.

## Handling Procedure

**Step 1** Log in to a Flume node as user **omm** and run the **ps -ef|grep "flume.role=server"** command to check whether the Flume process exists on the node.

- If yes, go to [Step 3](#).
- If no, restart the faulty Flume node or Flume service and go to [Step 2](#).

**Step 2** In the alarm list, check whether alarm "Flume Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

**Collect the fault information.**

**Step 3** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 4** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 5** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 6** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.196 ALM-24001 Flume Agent Exception

### Alarm Description

The Flume agent monitoring module monitors the Flume agent status. This alarm is generated when the Flume agent process is faulty (checked every 5 seconds) or the Flume agent fails to start (an alarm is reported immediately).

This alarm is cleared when the Flume agent process recovers, Flume agent starts successfully, and the alarm handling is completed.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24001	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
AgentId	Specifies the ID of the agent for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

The Flume agent instance for which the alarm is generated cannot provide services properly, and the data transmission tasks of the instance are temporarily interrupted. Real-time data is lost during real-time data transmission.

## Possible Causes

- The **JAVA\_HOME** directory does not exist, or the Java permission is incorrect.
- The Flume agent directory permission is incorrect.
- The Flume agent fails to start.

## Handling Procedure

**Check whether the JAVA\_HOME directory exists or whether the JAVA permission is correct.**

**Step 1** Log in to the host for which the alarm is generated as user **root**.

**Step 2** Obtain the installation directory of the Flume client for which the alarm is generated. (The value of **AgentId** can be obtained from **Location** of the alarm.)

```
ps -ef|grep AgentId | grep -v grep | awk -F 'conf-file ' '{print $2}' | awk -F 'fusioninsight' '{print $1}'
```



**Step 3** Run the `su - Flume installation user` command to switch to the Flume installation user and run the `cd Flume client installation directory/fusioninsight-flume-1.9.0/conf/` command to go to the Flume configuration directory.

**Step 4** Run the `cat ENV_VARS | grep JAVA_HOME` command.

**Step 5** Check whether the `JAVA_HOME` directory exists. If both the command output in [Step 4](#) and `ll $JAVA_HOME/` are not empty, the `JAVA_HOME` directory exists.

- If yes, go to [Step 7](#).
- If no, go to [Step 6](#).

**Step 6** Specify a correct `JAVA_HOME` directory, for example, `export JAVA_HOME=${BIGDATA_HOME}/common/runtime0/jdkVersion number`.

**Step 7** Run the `$JAVA_HOME/bin/java -version` command to check whether the Flume agent running user has the Java execution permission. If the Java version is displayed in the command output, the Java permission meets the requirement. Otherwise, the Java permission does not meet the requirement.

- If yes, go to [Step 9](#).
- If no, go to [Step 8](#).

 **NOTE**

`JAVA_HOME` is the environment variable exported during Flume client installation. You can also go to `Flume client installation directory/fusioninsight-flume-1.9.0/conf` and run the `cat ENV_VARS | grep JAVA_HOME` command to view the variable value.

**Step 8** Run the `chmod 750 $JAVA_HOME/bin/java` command to grant the Java execution permission to the Flume agent running user.

**Check the directory permission of the Flume agent.**

**Step 9** Log in to the host for which the alarm is generated as user `root`.

**Step 10** Run the following command to switch to the Flume agent installation directory:

```
cd Flume client installation directory/fusioninsight-flume-1.9.0/conf/
```

**Step 11** Run the `ls -al * -R` command to check whether any file owner is the user running the Flume agent.

- If yes, go to [Step 12](#).
- If no, run the `chown` command to change the file owner to the user who runs the Flume agent.

**Check the Flume agent configuration.**

**Step 12** Run the `cat properties.properties | grep spooldir` and `cat properties.properties | grep TAILDIR` commands to check whether the Flume source type is `spoolDir` or `tailDir`. If any command output is displayed, the Flume source type is `spoolDir` or `tailDir`.

- If yes, go to [Step 13](#).
- If no, go to [Step 17](#).

**Step 13** Check whether the data monitoring directory exists.

- If yes, go to [Step 15](#).

- If no, go to [Step 14](#).

 NOTE

Run the `cat properties.properties | grep spoolDir` command to view the spoolDir monitoring directory.

```
[root@fusioninsight-flume-1.9.0/conf]# cat properties.properties | grep spoolDir
client.sources.aal.spoolDir = /opt/liuxingcheng/flumeclient/sourcedata/flumesourcedata1
[root@fusioninsight-flume-1.9.0/conf]#
```

Run the `cat properties.properties | grep parentDir` command to view the tailDir monitoring directory.

```
[root@fusioninsight-flume-1.9.0/conf]# cat properties.properties | grep parentDir
server.sources.AAAA.filegroups.F1.parentDir = /tmp/flumetest/taildir_data
[root@fusioninsight-flume-1.9.0/conf]#
```

**Step 14** Specify a correct data monitoring directory.

**Step 15** Check whether the Flume agent user has the read, write, and execute permissions on the monitoring directory specified in [Step 13](#).

- If yes, go to [Step 17](#).
- If no, go to [Step 16](#).

 NOTE

Go to the monitoring directory as the Flume running user. If files can be created, the Flume running user has the read, write, and execute permissions on the monitoring directory.

**Step 16** Run the `chmod 777 Flume monitoring directory` command to grant the Flume agent running user the read, write, and execute permissions on the monitoring directory specified in [Step 13](#).

**Step 17** Check whether the components connected to the Flume sink are in safe mode.

- If yes, go to [Step 18](#).
- If no, go to [Step 23](#).

 NOTE

If the sinks in the `properties.properties` configuration file are the HDFS sink and HBase sink, and the configuration file contains a keytab file, the components connected to the Flume sink are in safe mode.

If the sink in the `properties.properties` configuration file is the Kafka sink and `*.security.protocol` is set to `SASL_PLAINTEXT` or `SASL_SSL`, Kafka connected to the Flume sink is in safe mode.

**Step 18** Run the `ll ketab path` command to check whether the keytab authentication path specified by the `*.kerberosKeytab` parameter in the configuration file exists.

- If yes, go to [Step 20](#).
- If no, go to [Step 19](#).

 NOTE

To view the keytab path, run the `cat properties.properties | grep keytab` command.

**Step 19** Change the value of `kerberosKeytab` in [Step 18](#) to the custom keytab path and go to [Step 21](#).

**Step 20** Go to [Step 18](#) and check whether the Flume agent running user has the permission to access the keytab authentication file. If the keytab path is returned, the user has the permission. Otherwise, the user does not have the permission.

- If yes, go to [Step 22](#).
- If no, go to [Step 21](#).

**Step 21** Run the **chmod 755 *ketab file*** command to grant the read permission on the keytab file specified in [Step 19](#), and restart the Flume process.

**Step 22** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 23](#).

#### Collect fault information.

**Step 23** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 24** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 25** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 26** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.197 ALM-24003 Flume Client Connection Interrupted

### Alarm Description

The alarm module monitors the port connection status on the Flume server. This alarm is generated if the Flume server fails to receive a connection message from the Flume client in three consecutive minutes.

This alarm is cleared after the Flume server receives a connection message from the Flume client.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24003	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
Client IP Address	Specifies the IP address of the Flume client.
Client Name	Specifies the agent name of the Flume client.
Sink Name	Specifies the sink name of Flume Agent.

## Impact on the System

The communication between the Flume client and the server fails. The Flume client cannot send data to the Flume server.

## Possible Causes

- The network connection between the Flume client and the server is faulty.
- The Flume client's process is abnormal.
- The Flume client is incorrectly configured.

## Handling Procedure

**Check the network connection between the Flume client and the server.**

**Step 1** Log in to the host whose IP address is specified by **Flume ClientIP** in the alarm information as user **root**.

**Step 2** Run the `ping Flume server IP address` command to check whether the network connection between the Flume client and the server is normal.

- If yes, go to [Step 3](#).
- If no, go to [Step 11](#).

**Check whether the Flume client's process is normal.**

**Step 3** Log in to the host whose IP address is specified by **Flume ClientIP** in the alarm information as user **root**.

**Step 4** Run the `ps -ef|grep flume |grep client` command to check whether the Flume client process exists.

- If yes, go to [Step 5](#).
- If no, go to [Step 11](#).

**Check the Flume client configuration.**

**Step 5** Log in to the host whose IP address is specified by **Flume ClientIP** in the alarm information as user **root**.

- Step 6** Run the `cd Flume client installation directory/fusioninsight-flume-1.9.0/conf/` command to go to Flume's configuration directory.
- Step 7** Run the `cat properties.properties` command to query the current configuration file of the Flume client.
- Step 8** Check whether the `properties.properties` file is correctly configured according to the configuration description of the Flume agent.
- If yes, go to [Step 9](#).
  - If no, go to [Step 11](#).
- Step 9** Modify the `properties.properties` configuration file.
- Check whether the alarm is cleared.**
- Step 10** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 11](#).
- Collect the fault information.**
- Step 11** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 12** Expand the **Service** drop-down list, and select **Flume** for the target cluster.
- Step 13** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 14** Collect logs in the `/var/log/Bigdata/flume-client` directory on the Flume client using a transmission tool.
- Step 15** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.198 ALM-24004 Exception Occurs When Flume Reads Data

## Alarm Description

The alarm module monitors the Flume source status. This alarm is generated when the duration in which the source cannot read data exceeds the threshold.

The default threshold is **0**, indicating that this function is disabled. You can change the threshold by modifying the `properties.properties` file in the `conf` directory. Specifically, modify the `NoDatatime` parameter of required the source.

The alarm is cleared when the source reads the data and the alarm handling is complete.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24004	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
AgentId	Specifies the ID of the agent for which the alarm was generated.
ComponentType	Specifies the type of the component for which the alarm was generated.
ComponentName	Specifies the name of the component for which the alarm was generated.

## Impact on the System

If data is found in the data source but the Flume source continuously fails to read data, the collection is stopped.

## Possible Causes

- The Flume source is faulty, so data cannot be sent.
- The network is faulty, so the data cannot be sent.

## Handling Procedure

**Check whether the Flume source is faulty.**

**Step 1** Open the **properties.properties** configuration file on the local PC, search for keyword **type = spooldir** in the file, and check whether the Flume source type is `spoolDir`.

- If yes, go to [Step 2](#).
- If no, go to [Step 3](#).

**Step 2** View the spoolDir monitoring directory to check whether all files are already transferred.

- If yes, no further action is required.
- If no, go to [Step 5](#).

 **NOTE**

The monitoring directory of spoolDir is specified by the `.spoolDir` parameter in the `properties.properties` configuration file. If all files in the monitoring directory have been transferred, the file name extension of all files in the monitoring directory is `.COMPLETED`.

**Step 3** Open the `properties.properties` configuration file on the local PC, search for `org.apache.flume.source.kafka.KafkaSource` in the file, and check whether the Flume source type is Kafka.

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

**Step 4** Check whether the topic data configured by the Kafka source has been used up.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Step 5** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Flume > Instances**.

**Step 6** Go to the Flume instance page of the faulty node to check whether the **Source Speed Metrics** in the alarm is **0**.

- If yes, go to [Step 11](#).
- If no, go to [Step 7](#).

**Check the network connectivity between the node with the IP address configured for the Flume source and the faulty node.**

**Step 7** Open the `properties.properties` configuration file on the local PC, search for `type = avro` in the file, and check whether the Flume source type is Avro.

- If yes, go to [Step 8](#).
- If no, go to [Step 11](#).

**Step 8** Log in to the faulty node as user `root`, and run the `ping IP address of the Flume source` command to check whether the peer host can be pinged successfully.

- If yes, go to [Step 11](#).
- If no, go to [Step 9](#).

**Step 9** Contact the network administrator to restore the network.

**Step 10** Wait for a while and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Collect fault information.**

**Step 11** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 12** Expand the **Service** drop-down list, and select **Flume** for the target cluster.
- Step 13** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 14** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.199 ALM-24005 Exception Occurs When Flume Transmits Data

### Alarm Description

The alarm module monitors the capacity status of Flume Channel. The alarm is generated immediately when the duration that Channel is fully occupied exceeds the threshold or the number of times that Source fails to send data to Channel exceeds the threshold.

The default threshold is **10**. You can change the threshold by modifying the **channelfullcount** parameter of the related channel in the **properties.properties** configuration file in the **conf** directory.

The alarm is cleared when the space of Flume Channel is released and the alarm handling is complete.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24005	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.



Parameter	Description
HostName	Specifies the host for which the alarm was generated.
AgentId	Specifies the ID of the agent for which the alarm was generated.
ComponentType	Specifies the type of the component for which the alarm was generated.
ComponentName	Specifies the name of the component for which the alarm was generated.

## Impact on the System

If the disk usage of Flume Channel increases continuously, the time required for importing data to a specified destination prolongs. When the disk usage of Flume Channel reaches 100%, the Flume agent process pauses.

## Possible Causes

- Flume Sink is faulty, so the data cannot be sent.
- The network is faulty, so the data cannot be sent.

## Handling Procedure

**Check whether Flume Sink is faulty.**

- Step 1** Open the **properties.properties** configuration file on the local PC, search for **type = hdfs** in the file, and check whether the Flume sink type is HDFS.
- If yes, go to [Step 2](#).
  - If no, go to [Step 3](#).
- Step 2** On FusionInsight Manager, check whether **HDFS Service Unavailable** alarm is generated in the alarm list and whether the HDFS service is stopped in the service list.
- If the alarm is reported, clear it according to the handling suggestions of ALM-14000 HDFS Service Unavailable; if the HDFS service is stopped, start it. Then, go to [Step 7](#).
  - If no, go to [Step 7](#).
- Step 3** Open the **properties.properties** configuration file on the local PC, search for **type = hbase** in the file, and check whether the Flume sink type is HBase.
- If yes, go to [Step 4](#).
  - If no, go to [Step 5](#).
- Step 4** On FusionInsight Manager, check whether **HBase Service Unavailable** alarm is generated in the alarm list and whether the HBase service is stopped in the service list.

- If the alarm is reported, clear it according to the handling suggestions of ALM-19000 HBase Service Unavailable; if the HBase service is stopped, start it. Then, go to [Step 7](#).
- If no, go to [Step 7](#).

**Step 5** Open the **properties.properties** configuration file on the local PC, search for **org.apache.flume.sink.kafka.KafkaSink** in the file, and check whether the Flume sink type is Kafka.

- If yes, go to [Step 6](#).
- If no, go to [Step 9](#).

**Step 6** On FusionInsight Manager, check whether **Kafka Service Unavailable** alarm is generated in the alarm list and whether the Kafka service is stopped in the service list.

- If the alarm is reported, clear it according to the handling suggestions of ALM-38000 Kafka Service Unavailable; if the Kafka service is stopped, start it. Then, go to [Step 7](#).
- If no, go to [Step 7](#).

**Step 7** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume** > **Instance**.

**Step 8** Go to the Flume instance page of the faulty node to check whether the indicator **Sink Speed Metrics** is 0.

- If yes, go to [Step 13](#).
- If no, go to [Step 9](#).

**Check the network connection between the faulty node and the node that corresponds to the Flume Sink IP address.**

**Step 9** Open the **properties.properties** configuration file on the local PC, search for **type = avro** in the file, and check whether the Flume sink type is Avro.

- If yes, go to [Step 10](#).
- If no, go to [Step 13](#).

**Step 10** Log in to the faulty node as user **root**, and run the **ping IP address of the Flume sink** command to check whether the peer host can be pinged successfully.

- If yes, go to [Step 13](#).
- If no, go to [Step 11](#).

**Step 11** Contact the network administrator to restore the network.

**Step 12** In the alarm list, check whether the alarm is cleared after a period.

- If yes, no further action is required.
- If no, go to [Step 13](#).

**Collect the fault information.**

**Step 13** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 14** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 15** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 16** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.200 ALM-24006 Heap Memory Usage of Flume Server Exceeds the Threshold

## Alarm Description

The system checks the heap memory usage of the Flume service every 60 seconds. This alarm is generated when the heap memory usage of the Flume instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24006	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the heap memory overflows, the service may break down and the Flume instance may be unavailable.

## Possible Causes

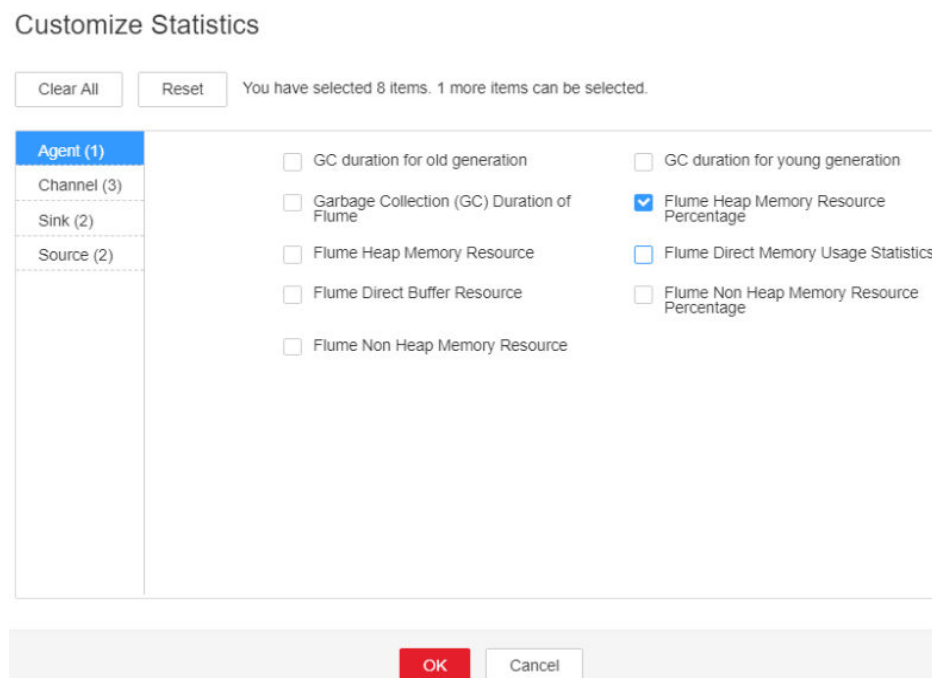
The heap memory of the Flume instance is overused or the heap memory is inappropriately allocated.

## Handling Procedure

**Check the heap memory usage.**

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **Flume Heap Memory Usage Exceeds the Threshold**, and view the **Location** information. Check the name of the host for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Name of the target cluster > Services > Flume**. On the page that is displayed, click the **Instance** tab. On the displayed tab page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Agent** and select **Flume Heap Memory Resource Percentage**. Then, click **OK**.

**Figure 7-120** Flume Heap Memory Resource Percentage



- Step 3** Check whether the heap memory used by Flume reaches the threshold (95% of the maximum heap memory by default).
- If yes, go to **Step 4**.

- If no, go to [Step 6](#).

**Step 4** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Flume** > **Configuration**. On the page that is displayed, click **All Configurations** and choose **Flume** > **System**. Set **-Xmx** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

 **NOTE**

If this alarm is generated, the heap memory configured for the Flume server is insufficient for data transmission. You are advised to change the heap memory to: Channel capacity x Maximum size of a single data record x Number of channels. Note that the value of **xmx** cannot exceed the remaining memory of the node.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

---

**NOTICE**

- During the restart, the Flume service is interrupted.
  - During the instance restart, if the failover mode of SinkGroup is configured and at least one instance is running properly, the Flume service is not interrupted. Otherwise, the Flume service is interrupted.
- 

**Collect the fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 7** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.201 ALM-24007 Flume Server Direct Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the direct memory usage of the Flume service every 60 seconds. This alarm is generated when the direct memory usage of the Flume instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the Flume direct memory usage is less than or equal to the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24007	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

If the direct memory overflows, the service may break down and the Flume instance may be unavailable.

### Possible Causes


The direct memory of the Flume process is overused or the direct memory is inappropriately allocated.

### Handling Procedure

**Check the direct memory usage.**

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **Flume Direct Memory Usage Exceeds the Threshold**, and view the **Location** information. Check the name of the host for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Name of the target cluster > Services > Flume**. On the page that is displayed, click the **Instance** tab. On the displayed tab page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Agent** and select **Flume Direct Memory Resource Percentage**. Then, click **OK**.

**Figure 7-121** Flume Direct Memory Usage Statistics  
Customize Statistics



- Step 3** Check whether the direct memory used by Flume reaches the threshold (80% of the maximum direct memory by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.
- Step 4** On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Service > Flume > Configuration**. On the page that is displayed, click **All Configurations** and choose **Flume > System**. Set **-XX:MaxDirectMemorySize** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

**NOTE**

If this alarm is generated, the direct memory size configured for the Flume server instance cannot meet service requirements. You are advised to change the value of **-XX:MaxDirectMemorySize** to twice the current direct memory size or change the value based on site requirements.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**NOTICE**

- During the restart, the Flume service is interrupted.
- During the instance restart, if the failover mode of SinkGroup is configured and at least one instance is running properly, the Flume service is not interrupted. Otherwise, the Flume service is interrupted.

**Collect the fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.202 ALM-24008 Flume Server Non Heap Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the non-heap memory usage of the Flume service every 60 seconds. This alarm is generated when the non-heap memory usage of the Flume instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24008	Major	Yes



## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the non-heap memory overflows, the service may break down and the Flume instance may be unavailable.

## Possible Causes


The non-heap memory of the Flume instance is overused or the non-heap memory is inappropriately allocated.

## Handling Procedure

**Check non-heap memory usage.**

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **Flume Non Heap Memory Usage Exceeds the Threshold**, and view the **Location** information. Check the name of the host for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Name of the target cluster > Services > Flume**. On the page that is displayed, click the **Instance** tab. On the displayed tab page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Agent** and select **Flume Non Heap Memory Resource Percentage**. Then, click **OK**.

**Figure 7-122** Flume Non-Heap Memory Resource Percentage



- Step 3** Check whether the non-heap memory used by Flume reaches the threshold (80% of the maximum non-heap memory by default).
- If yes, go to [Step 4](#).
  - If no, go to [Step 6](#).

- Step 4** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Flume** > **Configuration**. On the page that is displayed, click **All Configurations** and choose **Flume** > **System**. Set **-XX: MaxPermSize** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

**NOTE**

If this alarm is generated, the non-heap memory size configured for the Flume server instance cannot meet service requirements. You are advised to change the value of **-XX:MaxPermSize** to twice the current non-heap memory size or change the value based on site requirements.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).

**NOTICE**

- During the restart, the Flume service is interrupted.
- During the instance restart, if the failover mode of SinkGroup is configured and at least one instance is running properly, the Flume service is not interrupted. Otherwise, the Flume service is interrupted.

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Flume** for the target cluster.
- Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.203 ALM-24009 Flume Server Garbage Collection (GC) Time Exceeds the Threshold

### Alarm Description

The system checks the GC duration of the Flume process every 60 seconds. This alarm is generated when the GC duration of the Flume process exceeds the threshold (12 seconds by default) for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24009	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Parameter	Description
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Flume data transmission efficiency decreases.

## Possible Causes


The heap memory of the Flume process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

**Check the GC duration.**

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **GC Duration Exceeds the Threshold**, and view the **Location** information. Check the name of the host for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Name of the target cluster > Services > Flume**. On the page that is displayed, click the **Instance** tab. On the displayed tab page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Agent** and select **Garbage Collection (GC) Duration of Flume**. Then, click **OK**.

**Figure 7-123** Garbage Collection (GC) Duration of Flume



**Step 3** Check whether the GC duration of the Flume process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Flume** > **Configuration**. On the page that is displayed, click **All Configurations** and choose **Flume** > **System**. Set **-Xmx** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

 **NOTE**

If this alarm is generated, the heap memory configured for the Flume server is insufficient for data transmission. You are advised to change the heap memory to: Channel capacity x Maximum size of a single data record x Number of channels. Note that the value of **xmx** cannot exceed the remaining memory of the node.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

---

**NOTICE**

- During the restart, the Flume service is interrupted.
  - During the instance restart, if the failover mode of SinkGroup is configured and at least one instance is running properly, the Flume service is not interrupted. Otherwise, the Flume service is interrupted.
- 

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 7** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.204 ALM-24010 Flume Certificate File Is Invalid or Damaged

This section applies to MRS 3.2.0 or later.

### Alarm Description

Flume checks whether the Flume certificate file is valid (whether the certificate exists and whether the certificate format is correct) every hour. This alarm is generated when the certificate file is invalid or damaged. This alarm is automatically cleared when the certificate file becomes valid again.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24010	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System

The Flume client cannot access the Flume server.

### Possible Causes

The Flume certificate file is invalid or damaged.

### Handling Procedure

**View alarm information.**

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **ALM-24010 Flume Certificate File Is Invalid or Damaged**, and view

the **Location** information. View the IP address of the instance for which the alarm is generated.

**Check whether the certificate file in the system is valid. If it is not, generate a new one.**

**Step 2** Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

**Step 3** Run the following command to go to the Flume service certificate directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf
```

**Step 4** Run the **ls -l** command to check whether the **flume\_sChat.crt** file exists.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

**Step 5** Run the **openssl x509 -in flume\_sChat.crt -text -noout** command to check whether certificate details are displayed properly.

- If yes, go to **Step 9**.
- If no, go to **Step 6**.

**Step 6** Run the following command to go to the Flume script directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin
```

**Step 7** Run the following command to generate a new certificate file. Then check whether the alarm is automatically cleared one hour later.

```
sh geneJKS.sh -f Custom certificate password of the Flume role on the server -g Custom certificate password of the Flume role on the client
```

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

#### NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

**Step 8** Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 9**.
- If no, no further action is required.

**Collect fault information.**

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 10** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 11** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.205 ALM-24011 Flume Certificate File Is About to Expire

This section applies to MRS 3.2.0 or later.

## Alarm Description

Flume checks whether the Flume certificate file is about to expire every hour. This alarm is generated when the remaining validity period is at most 30 days. This alarm is automatically cleared when the remaining validity period is greater than 30 days.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24011	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.



## Impact on the System

Currently, there is no impact on the system.

## Possible Causes

The Flume certificate file is about to expire.

## Handling Procedure

**View alarm information.**

**Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **ALM-24011 Flume Certificate Is About to Expire**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

**Check whether the certificate file in the system is valid. If it is not, generate a new one.**

**Step 2** Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

**Step 3** Run the following command to go to the Flume service certificate directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf
```

**Step 4** Run the following command to check the effective time and expiration time of the Flume user certificate:

```
openssl x509 -noout -text -in flume_sChat.crt
```

**Step 5** Perform [Step 6](#) to [Step 7](#) during off-peak hours to update the certificate file as needed.

**Step 6** Run the following command to go to the Flume script directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin
```

**Step 7** Run the following command to generate a new certificate file. Then, check whether the alarm is automatically cleared one hour later.

```
sh geneJKS.sh -f Custom certificate password of the Flume role on the server -g Custom certificate password of the Flume role on the client
```

- If yes, go to [Step 9](#).
- If no, go to [Step 8](#).

### NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

**Step 8** Log in to the Flume node for which the alarm is generated as user **omm** and repeat **Step 6** to **Step 7**. Then, check whether the alarm is automatically cleared one hour later.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

**Step 9** Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 10**.
- If no, no further action is required.

**Collect fault information.**

**Step 10** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 11** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 12** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.206 ALM-24012 Flume Certificate File Has Expired

This section applies to MRS 3.2.0 or later.

### Alarm Description

Flume checks whether its certificate file in the system has expired every hour. This alarm is generated when the server certificate has expired. This alarm is automatically cleared when the certificate file becomes valid again.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24012	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

The Flume client cannot access the Flume server.

## Possible Causes

The Flume certificate file has expired.

## Handling Procedure

**View alarm information.**

**Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **ALM-24012 Flume Certificate Has Expired**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

**Check whether the certificate file in the system is valid. If it is not, generate a new one.**

**Step 2** Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

**Step 3** Run the following command to go to the Flume service certificate directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf
```

**Step 4** Run the following command to check the effective time and expiration time of the HA user certificate to determine whether the certificate file is still in the validity period:

```
openssl x509 -noout -text -in flume_sChat.crt
```

- If yes, go to [Step 9](#).
- If no, go to [Step 5](#).

**Step 5** Run the following command to go to the Flume script directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/
flume/bin
```

**Step 6** Run the following command to generate a new certificate file. Then, check whether the alarm is automatically cleared one hour later.

```
sh geneJKS.sh -f Custom certificate password of the Flume role on the server -g
Custom certificate password of the Flume role on the client
```

- If yes, go to [Step 8](#).
- If no, go to [Step 7](#).

 **NOTE**

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

**Step 7** Log in to the Flume node for which the alarm is generated as user **omm** and repeat [Step 5](#) to [Step 6](#). Then, check whether the alarm is automatically cleared one hour later.

- If yes, go to [Step 8](#).
- If no, go to [Step 9](#).

**Step 8** Check whether this alarm is generated again during periodic system check.

- If yes, go to [Step 9](#).
- If no, no further action is required.

**Collect fault information.**

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Expand the **Service** drop-down list, and select **Flume** for the target cluster.

**Step 11** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.207 ALM-24013 Flume MonitorServer Certificate File Is Invalid or Damaged

This section applies to MRS 3.2.0 or later.

### Alarm Description

MonitorServer checks whether its certificate file is valid (whether the certificate exists and whether the certificate format is correct) every hour. This alarm is generated when the certificate file is invalid or damaged. This alarm is automatically cleared when the certificate file becomes valid again.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24013	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System

The Flume client cannot access the Flume server.

### Possible Causes

The MonitorServer certificate file is invalid or damaged.

### Handling Procedure

**View alarm information.**

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **ALM-24013 MonitorServer Certificate File Is Invalid or Damaged**,

and view the **Location** information. View the IP address of the instance for which the alarm is generated.

**Check whether the certificate file in the system is valid. If it is not, generate a new one.**

**Step 2** Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

**Step 3** Run the following command to go to the MonitorServer certificate file directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf
```

**Step 4** Run the **ls -l** command to check whether the **ms\_sChat.crt** file exists:

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

**Step 5** Run the **openssl x509 -in ms\_sChat.crt -text -noout** command to check whether certificate details are displayed.

- If yes, go to **Step 9**.
- If no, go to **Step 6**.

**Step 6** Run the following command to go to the Flume script directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin
```

**Step 7** Run the following command to generate a new certificate file. Then check whether the alarm is automatically cleared one hour later.

```
sh geneJKS.sh -m Custom password of the MonitorServer certificate on the server
-n Custom password of the MonitorServer certificate on the client
```

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

#### NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

**Step 8** Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 9**.
- If no, no further action is required.

**Collect fault information.**

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 10** Select **MonitorServer** in the required cluster for **Service**.

**Step 11** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.208 ALM-24014 Flume MonitorServer Certificate Is About to Expire

This section applies to MRS 3.2.0 or later.

### Alarm Description

MonitorServer checks whether its certificate file is about to expire every hour. This alarm is generated when the remaining validity period is at most 30 days. This alarm is automatically cleared when the remaining validity period is greater than 30 days.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24014	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

Currently, there is no impact on the system.

## Possible Causes

The MonitorServer certificate file is about to expire.

## Handling Procedure

**View alarm information.**

**Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **ALM-24014 MonitorServer Certificate Is About to Expire**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

**Check whether the certificate file in the system is valid. If it is not, generate a new one.**

**Step 2** Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

**Step 3** Run the following command to go to the MonitorServer certificate file directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf
```

**Step 4** Run the following command to check the effective time and expiration time of the MonitorServer user certificate:

```
openssl x509 -noout -text -in ms_sChat.crt
```

**Step 5** Perform [Step 6](#) to [Step 7](#) during off-peak hours to update the certificate file as needed.

**Step 6** Run the following command to go to the Flume script directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin
```

**Step 7** Run the following command to generate a new certificate file. Then, check whether the alarm is automatically cleared one hour later.

```
sh geneJKS.sh -m Custom password of the MonitorServer certificate on the server
-n Custom password of the MonitorServer certificate on the client
```

- If yes, go to [Step 9](#).
- If no, go to [Step 8](#).

### NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.



**Step 8** Log in to the Flume node for which the alarm is generated as user **omm** and repeat **Step 6** to **Step 7**. Then, check whether the alarm is automatically cleared one hour later.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

**Step 9** Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 10**.
- If no, no further action is required.

**Collect fault information.**

**Step 10** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 11** Select **MonitorServer** in the required cluster for **Service**.

**Step 12** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.209 ALM-24015 Flume MonitorServer Certificate File Has Expired

This section applies to MRS 3.2.0 or later.

## Alarm Description

MonitorServer checks whether its certificate file in the system has expired every hour. This alarm is generated when the server certificate has expired. This alarm is automatically cleared when the MonitorServer certificate file becomes valid again.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24015	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

The Flume client cannot access the Flume server.

## Possible Causes

The MonitorServer certificate file has expired.

## Handling Procedure

**View alarm information.**

**Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing **ALM-24015 MonitorServer Certificate Has Expired**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

**Check whether the certificate file in the system is valid. If it is not, generate a new one.**

**Step 2** Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

**Step 3** Run the following command to go to the MonitorServer certificate file directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf
```

**Step 4** Run the following command to check the effective time and expiration time of the user certificate to determine whether the certificate file is still in the validity period:

```
openssl x509 -noout -text -in ms_sChat.crt
```

- If yes, go to [Step 9](#).
- If no, go to [Step 5](#).

**Step 5** Run the following command to go to the Flume script directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/
flume/bin
```

**Step 6** Run the following command to generate a new certificate file. Then, check whether the alarm is automatically cleared one hour later.

```
sh geneJKS.sh -m Custom password of the MonitorServer certificate on the server
-n Custom password of the MonitorServer certificate on the client
```

- If yes, go to [Step 8](#).
- If no, go to [Step 7](#).

#### NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

**Step 7** Log in to the Flume node for which the alarm is generated as user **omm** and repeat [Step 5](#) to [Step 6](#). Then, check whether the alarm is automatically cleared one hour later.

- If yes, go to [Step 8](#).
- If no, go to [Step 9](#).

**Step 8** Check whether this alarm is generated again during periodic system check.

- If yes, go to [Step 9](#).
- If no, no further action is required.

#### Collect fault information.

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Select **MonitorServer** in the required cluster for **Service**.

**Step 11** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.210 ALM-25000 LdapServer Service Unavailable

### Description

The system checks the LdapServer service status every 30 seconds. This alarm is generated when the system detects that both the active and standby LdapServer services are abnormal.

This alarm is cleared when the system detects that one or two LdapServer services are normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
25000	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System


The running status of the component that depends on the LdapServer becomes faulty. As a result, Kerberos authentication fails in the cluster or OS user cache synchronization is abnormal, and component services are abnormal.

### Possible Causes

- The node where the LdapServer service locates is faulty.
- The LdapServer process is abnormal.

### Procedure

**Check whether the nodes where the two SlapdServer instances of the LdapServer service are located are faulty.**

- Step 1** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **LdapServer** > **Instance** to go to the LdapServer instance page to obtain the host name of the node where the two SlapdServer instances locates.
- Step 2** Choose **O&M** > **Alarm** > **Alarms**. On the **Alarm** page of the FusionInsight Manager system, check whether any alarm of **Node Fault** exists.
- If yes, go to **Step 3**.
  - If no, go to **Step 6**.
- Step 3** Check whether the host name in the alarm is consistent with the **Step 1** host name.
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.
- Step 4** Handle the alarm according to "ALM-12006 Node Fault".
- Step 5** Check whether **LdapServer Service Unavailable** is cleared in the alarm list.
- If yes, no further action is required.
  - If no, go to **Step 10**.
- Check whether the LdapServer process is normal.**
- Step 6** Choose **O&M** > **Alarm** > **Alarms**. On the **Alarm** page of the FusionInsight Manager system, check whether any alarm of **Process Fault** exists.
- If yes, go to **Step 7**.
  - If no, go to **Step 10**.
- Step 7** Check whether the service and host name in the alarm are consistent with the LdapServer service and host name.
- If yes, go to **Step 8**.
  - If no, go to **Step 10**.
- Step 8** Handle the alarm according to "ALM-12007 Process Fault".
- Step 9** Check whether **LdapServer Service Unavailable** is cleared in the alarm list.
- If yes, no further action is required.
  - If no, go to **Step 10**.
- Collect fault information.**
- Step 10** On the FusionInsight Manager, choose **O&M** > **Log** > **Download**.
- Step 11** Select **LdapServer** in the required cluster from the **Service**.
- Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 13** Contact the O&M personnel and send the collected logs.
- End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.211 ALM-25004 Abnormal LdapServer Data Synchronization

### Description

The system checks the LdapServer data every 30 seconds. This alarm is generated when the data on the active and standby LdapServers of Manager is inconsistent for 12 consecutive times. This alarm is cleared when the data on the active and standby LdapServers is consistent.

The system checks the LdapServer data every 30 seconds. This alarm is generated when the LdapServer data in the cluster is inconsistent with that on Manager for 12 consecutive times. This alarm is cleared when the data is consistent.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
25004	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

LdapServer data inconsistency occurs because the LdapServer data in Manager is damaged or the LdapServer data in the cluster is damaged. The LdapServer process with damaged data cannot provide services externally, and the authentication functions of Manager and the cluster are affected.

### Possible Causes

- The network of the node where the LdapServer process locates is faulty.

- The LdapServer process is abnormal.
- The OS restart damages data on LdapServer.
- The amount of Oldap data exceeds the threshold (10 MB by default).

## Procedure

### Check whether the network where the LdapServer nodes reside is faulty.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. Record the IP address of HostName in the alarm locating information as IP1 (if multiple alarms exist, record the IP addresses as IP1, IP2, and IP3 respectively).
- Step 2** Contact O&M personnel and log in to the nodes corresponding to IP 1. Run the ping command to check whether the IP address of the management plane of the active OMS node can be pinged.
- If yes, go to [Step 4](#).
  - If no, go to [Step 3](#).
- Step 3** Contact the network administrator to recover the network and check whether **Abnormal LdapServer Data Synchronization** is cleared.
- If yes, no further action is required.
  - If no, go to [Step 4](#).

### Check whether the LdapServer processes are normal.

- Step 4** On the **Alarm** page of FusionInsight Manager, check whether the **OLdap Resource Abnormal** exists.
- If yes, go to [Step 5](#).
  - If no, go to [Step 7](#).
- Step 5** Clear the alarm by following the steps provided in "ALM-12004 OLdap Resource Abnormal".
- Step 6** Check whether **Abnormal LdapServer Data Synchronization** is cleared in the alarm list.
- If yes, no further action is required.
  - If no, go to [Step 7](#).
- Step 7** On the **Alarm** page of FusionInsight Manager, check whether **Process Fault** is generated for the LdapServer service.
- If yes, go to [Step 8](#).
  - If no, go to [Step 10](#).
- Step 8** Handle the alarm according to "ALM-12007 Process Fault".
- Step 9** Check whether **Abnormal LdapServer Data Synchronization** is cleared.
- If yes, no further action is required.
  - If no, go to [Step 10](#).

### Check whether the LdapServer processes are normal.

- Step 10** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Record the IP address of HostName in the alarm locating information as "IP1" (if multiple alarms exist,

record the IP addresses as "IP1", "IP2", and "IP3" respectively). Choose **Cluster > Name of the desired cluster > Services > LdapServer > Configurations**. Record the port number of LdapServer as "PORT". (If the IP address in the alarm locating information is the IP address of the standby management node, choose **System > OMS > oldap > Modify Configuration** and record the listening port number of LdapServer.)

**Step 11** Log in to the nodes corresponding to IP1 as user **omm**.

**Step 12** Run the following command to check whether errors are displayed in the queried information.

```
ldapsearch -H ldaps://IP1:PORT -LLL -x -D cn=root,dc=hadoop,dc=com -W -b ou=Peoples,dc=hadoop,dc=com
```

After running the command, enter the **LDAP** administrator password. Contact the system administrator to obtain the password.

- If yes, go to [Step 13](#).
- If no, go to [Step 15](#).

**Step 13** Recover the LdapServer and OMS nodes using data backed up before the alarm is generated.

 **NOTE**

Use the OMS data and LdapServer data backed up at the same point in time to recover the data. Otherwise, the service and operation may fail. To recover data when services run properly, you are advised to manually back up the latest management data and then recover the data. Otherwise, Manager data produced between the backup point in time and the recovery point in time will be lost.

**Step 14** Check whether alarm **Abnormal LdapServer Data Synchronization** is cleared.

- If yes, no further action is required.
- If no, go to [Step 15](#).

**Check whether the data volume of the Oldap exceeds the threshold** (10 MB by default). (This step applies only to versions earlier than MRS 3.3.0. For MRS 3.3.0 and later versions, go to [Step 18](#).)

**Step 15** Log in to the active OMS node as user **omm**.

**Step 16** Run the following command to check whether the directory contains **.mdb** files.

```
ll /srv/BigData/ldapData/oldap/data/
```

- If yes, check and record the size of the **.mdb** file and go to [Step 17](#).
- If no, go to [Step 18](#).

**Step 17** Run the following command to view the Oldap configuration and record the value of **Map size** (the default value is **10485760** bytes, that is, 10 MB)

```
mdb_stat -e /srv/BigData/ldapData/oldap/data/
```

Check whether the size of the **.mdb** file with [Step 16](#) records reaches the value of **Map size**.

- If yes, contact the O&M personnel.




- If no, go to [Step 18](#).

**Collect fault information.**

**Step 18** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 19** Select **LdapServer** in the required cluster and **OmsLdapServer** from the **Service**.

**Step 20** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 21** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.212 ALM-25005 nscd Service Exception

### Alarm Description

The system checks the status of the nscd service every 60 seconds. This alarm is generated when the nscd process fails to be queried for four consecutive times (three minutes) or users in LdapServer cannot be obtained.

This alarm is cleared when the process is restored and users in LdapServer can be obtained.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
25005	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

The alarmed node may not be able to synchronize data from LdapServer. The **id** command may fail to obtain the LDAP data, affecting upper-layer services.

## Possible Causes

- The nscd service is not started.
- The network is faulty, and cannot access the LDAP server.
- NameService is abnormal.
- Users cannot be queried because the OS executes commands too slowly.

## Handling Procedure

**Check whether the nscd service is started.**

**Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Record the IP address of **HostName** in **Location** of the alarm as **IP1** (if multiple alarms exist, record the IP addresses as **IP1**, **IP2**, and **IP3** respectively).

**Step 2** Contact the O&M personnel to access the node using IP1 as user **root**. Run the **ps -ef | grep nscd** command on the node and check whether the **/usr/sbin/nscd** process is started.

- If yes, go to [Step 5](#).
- If no, go to [Step 3](#).

**Step 3** Run the **service nscd restart** command as user **root** to restart the nscd service. Then run the **ps -ef | grep nscd** command to check whether the nscd service is started.

- If yes, go to [Step 4](#).
- If no, go to [Step 15](#).

**Step 4** Wait for 5 minutes and run the **ps -ef | grep nscd** command again as user **root**. Check whether the service exists.

- If yes, go to [Step 11](#).
- If no, go to [Step 15](#).

**Check whether the network is faulty, and whether the LDAP server can be accessed.**

**Step 5** Log in to the alarmed node as user **root** and run the **ping** command to check whether the network connectivity between this node and the LdapServer node is normal.

- If yes, go to [Step 6](#).
- If no, contact network administrators to troubleshoot the fault.

**Check whether the NameService is normal.**

**Step 6** Log in to the alarmed node as user **root**. Run the **cat /etc/nsswitch.conf** command to check whether the **passwd**, **group**, **services**, **netgroup**, and **aliases** of NameService are correctly configured.

The correct parameter configurations are as follows:

**passwd: compat ldap; group: compat ldap; services: files ldap; netgroup: files ldap; aliases: files ldap**

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

**Step 7** Log in to the alarmed node as user **root**. Run the **cat /etc/nscd.conf** command to check whether the **enable-cache passwd**, **positive-time-to-live passwd**, **enable-cache group**, and **positive-time-to-live group** in the configuration file are correctly configured.

The correct parameter configurations are as follows:

**enable-cache passwd: yes; positive-time-to-live passwd: 600; enable-cache group: yes; positive-time-to-live group: 3600**

- If yes, go to [Step 8](#).
- If no, go to [Step 10](#).

**Step 8** Run the **/usr/sbin/nscd -i group** and **/usr/sbin/nscd -i passwd** commands as user **root**. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If yes, go to [Step 11](#).
- If no, go to [Step 15](#).

**Step 9** Run the **vi /etc/nsswitch.conf** command as user **root**. Correct the configurations in [Step 6](#) and save the file. Run the **service nscd restart** command to restart the nscd service. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If yes, go to [Step 11](#).
- If no, go to [Step 15](#).

**Step 10** Run the **vi /etc/nscd.conf** command as user **root**. Correct the configurations in [Step 7](#) and save the file. Run the **service nscd restart** command to restart the nscd service. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If yes, go to [Step 11](#).
- If no, go to [Step 15](#).

**Step 11** Log in to the FusionInsight Manager portal. Wait for 5 minutes and check whether the **nscd Service Exception** alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

**Check whether frame freezing occurs when running a command in the operating system.**

**Step 12** Log in to the faulty node as user **root**, run the **id admin** command, and check whether the command execution takes a long time. If the command execution takes more than 3 seconds, the command execution is deemed to be slow.

- If yes, go to [Step 13](#).

- If no, go to [Step 15](#).

**Step 13** Run the `cat /var/log/messages` command to check whether the `nscd` frequently restarts or the error information "Can't contact LDAP server" exists.

`nscd` exception example:

```
Feb 11 11:44:42 10-120-205-33 nscd: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.55:21780:
Can't contact LDAP server
Feb 11 11:44:43 10-120-205-33 ntpq: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.55:21780:
Can't contact LDAP server
Feb 11 11:44:44 10-120-205-33 ntpq: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.92:21780:
Can't contact LDAP server
```

- If yes, go to [Step 14](#).
- If no, go to [Step 15](#).

**Step 14** Run the `vi$BIGDATA_HOME/tmp/random_ldap_ip_order` command to modify the number at the end. If the original number is an odd number, change it to an even number. If the number is an even number, change it to an odd number.

Run the `vi /etc/ldap.conf` command to enter the editing mode, press **Insert** to start editing, and then change the first two IP addresses of the URI configuration item.

After the modification is complete, press **Esc** to exit the editing mode and enter `:wq!` to save the settings and exit.


Run the `service nscd restart` command to restart the `nscd` service. Wait 5 minutes and run the `id admin` command again. Check whether the command execution is slow.

- If yes, go to [Step 15](#).
- If no, log in to other faulty nodes and repeat [Step 12](#) to [Step 14](#) to check whether the first `LdapServer` node in the URI before modifying `/etc/ldap.conf` is faulty. For example, check whether the service IP address is unreachable, the network delay is too long, or other abnormal software is deployed.

**Collect the fault information.**

**Step 15** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 16** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **LdapClient** for the target cluster.

**Step 17** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 18** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.213 ALM-25006 Sssd Service Exception

### Description

The system checks the status of the sssd service every 60 seconds. This alarm is generated when the sssd process fails to be queried for four consecutive times (three minutes) or users in LdapServer cannot be obtained.

This alarm is cleared when the process is restored and users in LdapServer can be obtained.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
25006	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.

### Impact on the System

The alarmed node may not be able to synchronize data from LdapServer. The id command may fail to obtain the LDAP data, affecting upper-layer services.

### Possible Causes

- The sssd service is not started or is incorrectly started.
- The network is faulty and cannot access the LDAP server.
- NameService is abnormal.
- Users cannot be queried because the OS executes commands too slowly.

### Procedure

**Check whether the sssd service is correctly started.**

**Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. Find the IP address of **HostName** in **Location** of the alarm and record it as IP1 (if multiple alarms exist, record the IP addresses as IP1, IP2, and IP3 respectively).

**Step 2** Contact the O&M personnel to access the node using IP1 as user **root**. Run the **ps -ef | grep sssd** command and check whether the **/usr/sbin/sss**d process is started.

- If the process is started, go to **Step 3**.
- If the process is not started, go to **Step 4**.

**Step 3** Check whether the sssd process queried in **Step 2** has three subprocesses.

- If yes, go to **Step 5**.
- If no, go to **Step 4**.

**Step 4** Run the **service sssd restart** command as user **root** to restart the sssd service. Then run the **ps -ef | grep sssd** command to check whether the sssd process is normal.

In the normal state, the **/usr/sbin/sss**d process has three subprocesses: **/usr/libexec/sss/sss\_be**, **/usr/libexec/sss/sss\_nss**, and **/usr/libexec/sss/sss\_pam**.

- If it exists, go to **Step 9**.
- If it does not exist, go to **Step 13**.

**Check whether the LDAP server can be accessed.**

**Step 5** Log in to the alarmed node as user **root**. Run the **ping** command to check the network connectivity between this node and the LdapServer node.

- If the network is normal, go to **Step 6**.
- If the network is faulty, contact network administrators to troubleshoot the fault.

**Check whether NameService is normal.**

**Step 6** Log in to the alarmed node as user **root**. Run the **cat /etc/nsswitch.conf** command and check the **passwd** and **group** configurations of NameService.

The correct parameter configurations are as follows: **passwd: compat ldap** and **group: compat ldap**.

- If the configurations are correct, go to **Step 7**.
- If the configurations are incorrect, go to **Step 8**.

**Step 7** Run the **/usr/sbin/sss\_cache -G** and **/usr/sbin/sss\_cache -U** commands as user **root**. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If results are queried, go to **Step 9**.
- If no result is queried, go to **Step 13**.

**Step 8** Run the **vi /etc/nsswitch.conf** command as user **root**. Correct the configurations in **Step 6** and save the file. Run the **service sssd restart** command to restart the sssd service. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If results are queried, go to **Step 9**.

- If no result is queried, go to [Step 13](#).

**Step 9** Log in to the FusionInsight Manager portal. Wait for 5 minutes and check whether the **sssd Service Exception** alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm persists, go to [Step 10](#).

**Check whether frame freezing occurs when running a command in the operating system.**

**Step 10** Log in to the faulty node as user **root**, run the **id admin** command, and check whether the command execution takes a long time. If the command execution takes more than 3 seconds, the command execution is deemed to be slow.

- If yes, go to [Step 11](#).
- If no, go to [Step 13](#).

**Step 11** Run the **cat /var/log/messages** command to check whether the sssd frequently restarts or the error information **Can't contact LDAP server** exists.

sssd restart example:

```
Feb 7 11:38:16 10-132-190-105 sssd[pam]: Shutting down
Feb 7 11:38:16 10-132-190-105 sssd[nss]: Shutting down
Feb 7 11:38:16 10-132-190-105 sssd[nss]: Shutting down
Feb 7 11:38:16 10-132-190-105 sssd[be[default]]: Shutting down
Feb 7 11:38:16 10-132-190-105 sssd: Starting up
Feb 7 11:38:16 10-132-190-105 sssd[be[default]]: Starting up
Feb 7 11:38:16 10-132-190-105 sssd[nss]: Starting up
Feb 7 11:38:16 10-132-190-105 sssd[pam]: Starting up
```

- If yes, go to [Step 12](#).
- If no, go to [Step 13](#).

**Step 12** Run the **vi \$BIGDATA\_HOME/tmp/random\_ldap\_ip\_order** command to modify the number at the end. If the original number is an odd number, change it to an even number. If the number is an even number, change it to an odd number.

Run the **vi /etc/sss/sss.conf** command to reverse the first two IP addresses of the **ldap\_uri** configuration item, save the settings, and exit.

Run the **ps -ef | grep sssd** command to query the ID of the sssd process, kill it, and run the **/usr/sbin/sss -D -f** command to restart the sssd service. Wait 5 minutes and run the **id admin** command again.


Check whether the command execution is slow.

- If yes, go to [Step 13](#).
- If no, log in to other faulty nodes and run [Step 10](#) to [Step 12](#). Collect logs and check whether the first ldapserver node in the ldap\_uri before modifying **/etc/sss/sss.conf** is faulty. For example, check whether the service IP address is unreachable, the network latency is too long, or other abnormal software is deployed.

**Collect fault information.**

**Step 13** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 14** Select **LdapClient** in the required cluster from the **Service**.

**Step 15** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 16** Contact the O&M personnel and send the collected fault logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.214 ALM-25007 Number of SlapdServer Connections Exceeds the Threshold

### Alarm Description

The system checks the number of process connections on the SlapdServer node every 30 seconds and compares the actual number with the threshold. This alarm is generated when the number of process connections exceeds the threshold (1000 by default) for multiple times (5 by default).

Its **Trigger Count** is configurable. If **Trigger Count** is set to 1, this alarm is cleared when the number of process connections is less than or equal to the threshold. If **Trigger Count** is greater than 1, this alarm is cleared when the number of process connections is less than or equal to 90% of the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
25007	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.



Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

SlapdServer may respond slowly or become unavailable. Kerberos authentication times out or OS user caches cannot be synchronized. Component services are faulty.

## Possible Causes

- There are too many SlapdServer connections.
- The alarm threshold or alarm trigger count is improperly configured.

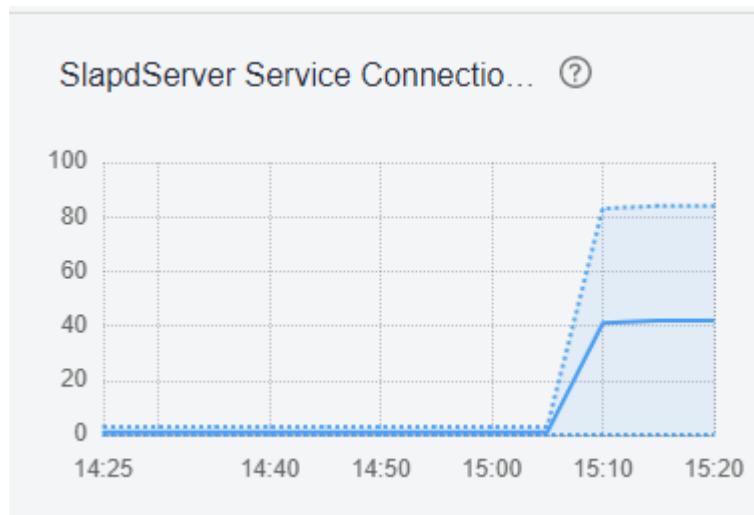
## Handling Procedure

**Check whether there are too many SlapdServer process connections.**

**Step 1** Log in to FusionInsight Manager and choose **Cluster > Services > LdapServer**.

**Step 2** On the LdapServer dashboard page, observe the SlapdServer process connections and decrease the connections based on service requirements.


**Figure 7-124** SlapdServer process connections



**Step 3** Wait about 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the alarm threshold or alarm trigger count is properly configured.**

- Step 4** On FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, choose **LdapServer > Other > SlapdServer Service Connections**, and check whether the alarm trigger count and alarm threshold are set properly.
- If yes, go to **Step 7**.
  - If no, go to **Step 5**.
- Step 5** Change the trigger count and alarm threshold based on the actual number of process connections, and apply the changes.
- Step 6** Wait 2 minutes and check whether the alarm is automatically cleared.
- If yes, no further action is required.
  - If no, go to **Step 7**.
- Collect fault information.**
- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 8** Expand the **Service** drop-down list, and select **LdapServer** for the target cluster.
- Step 9** Specify **Hosts** for collecting logs, which is optional. By default, all hosts are selected.
- Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 11** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.215 ALM-25008 SlapdServer CPU Usage Exceeds the Threshold

### Alarm Description

The system checks the CPU usage of the SlapdServer node every 30 seconds and compares the actual usage with the threshold. This alarm is generated when the SlapdServer CPU usage exceeds the threshold for multiple times (5 by default).

Its **Trigger Count** is configurable. If **Trigger Count** is set to **1**, this alarm is cleared when the SlapdServer CPU usage is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the SlapdServer CPU usage is less than or equal to 90% of the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
25008	Critical (default threshold: 85%) Major (default threshold: 75%)	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

SlapdServer may respond slowly or become unavailable. Kerberos authentication times out or OS user caches cannot be synchronized. Component services are faulty.

## Possible Causes

- The alarm threshold or alarm trigger count is improperly configured.
- The CPU configuration cannot meet service requirements, and the CPU usage reaches the upper limit.

## Handling Procedure

**Check whether the alarm threshold or alarm trigger count is properly configured.**

**Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, choose **LdapServer > Other > SlapdServer Service Total CPU Percentage**, and check whether the alarm trigger count and alarm threshold are set properly.

- If yes, go to **Step 4**.

- If no, go to [Step 2](#).

**Step 2** Change the trigger count and alarm threshold based on the actual CPU usage, and apply the changes.

**Step 3** Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the CPU usage reaches the upper limit.**

**Step 4** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the right pane, click this alarm and obtain the host name in **Location**.

**Step 5** Choose **Cluster > Services > LdapServer**, click the **Instance** tab, and click the SlapdServer instance corresponding to the host name in [Step 4](#).

**Step 6** On the dashboard of the instance, observe the real-time data of the **CPU Usage of a Single SlapdServer Instance** chart for about 5 minutes and check whether the CPU usage exceeds the threshold (**75%** by default) for multiple times.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

**Step 7** Check whether the status of other SlapdServer instances is normal. For details, see [Step 5](#) to [Step 6](#).

- If yes, contact the MRS cluster administrator to evaluate whether to expand the capacity of SlapdServer instances. Then, go to [Step 8](#).
- If no, repair the faulty SlapdServer instance and go to [Step 8](#).


**Step 8** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Collect fault information.**

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Expand the **Service** drop-down list, and select **LdapServer** for the target cluster.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.216 ALM-25500 KrbServer Service Unavailable

### Description

The system checks the KrbServer service status every 30 seconds. This alarm is generated when the system detects that the KrbServer service is abnormal.

This alarm is cleared when the system detects that the KrbServer service is normal.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
25500	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

The running status of the component that depends on the KrbServer becomes faulty. As a result, the Kerberos authentication of the cluster fails, and the component services are abnormal.

### Possible Causes

- The node where the KrbServer service locates is faulty.
- The OLdap service is abnormal.

### Procedure

**Check whether the node where the KrbServer service locates is faulty.**

- Step 1** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **KrbServer** > **Instance** to go to the KrbServer instance page to obtain the host name of the node where the KrbServer service locates.

**Step 2** On the **Alarm** page of the FusionInsight Manager system, check whether any alarm of **Node Fault** exists.

- If yes, go to [Step 3](#).
- If no, go to [Step 6](#).

**Step 3** Check whether the host name in the alarm is consistent with the [Step 1](#) host name.

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** Handle the alarm according to "ALM-12006 Node Fault".

**Step 5** Check whether **KrbServer Service Unavailable** is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Check whether the OLdap service is normal.**

**Step 6** On the **Alarm** page of the FusionInsight Manager system, check whether any alarm of **OLdap Resource Abnormal** exists.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

**Step 7** Handle the alarm according to "ALM-12004 OLdap Resource Abnormal".


**Step 8** Check whether **KrbServer Service Unavailable** is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Collect fault information.**

**Step 9** On the FusionInsight Manager, choose **O&M > Log > Download**.

**Step 10** Select **KrbServer** in the required cluster from the **Service**.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.217 ALM-26051 Storm Service Unavailable

### Description

The system checks the Storm service status every 30 seconds. This alarm is generated when all Nimbus nodes in the cluster are abnormal and the Storm service is unavailable.

This alarm is cleared when the Storm service recovers.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
26051	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

The cluster cannot provide the Storm service, and users cannot perform new Storm tasks.

### Possible Causes

- The Kerberos cluster is faulty.
- The ZooKeeper cluster is faulty or suspended.
- The active and standby Nimbus nodes in the Storm cluster are abnormal

### Procedure

**Check the status of the Kerberos cluster. (Skip this step if the normal mode is used.)**

- Step 1** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services**.

- Step 2** Check whether the running status of the Kerberos service is **Normal**.
- If yes, go to [Step 5](#).
  - If no, go to [Step 3](#).

**Step 3** See the related maintenance information of **ALM-25500 KrbServer Service Unavailable**.

- Step 4** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 5](#).

#### **Check the status of the ZooKeeper cluster.**

- Step 5** Check whether the running status of the ZooKeeper service is **Normal**.
- If yes, go to [Step 8](#).
  - If no, go to [Step 6](#).

**Step 6** If ZooKeeper service is stopped, start it, else see the related maintenance information of **ALM-13000 ZooKeeper Service Unavailable**.

- Step 7** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 8](#).

#### **Check the status of the active and standby Nimbus nodes.**

**Step 8** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Nimbus** to go to the Nimbus Instances page.

- Step 9** Check whether only one Nimbus node that is in the **Active** state in **Roles**.
- If yes, go to [Step 13](#).
  - If no, go to [Step 10](#).

- Step 10** Select two Nimbus role instances, choose **More** > **Restart Instance**, and check whether the instances restart successfully.
- If yes, go to [Step 11](#).
  - If no, go to [Step 13](#).

- Step 11** Log in to the FusionInsight Manager portal again, choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Nimbus** to check whether the running status is **Normal**.
- If yes, go to [Step 12](#).
  - If no, go to [Step 13](#).

- Step 12** Wait for 30 seconds and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 13](#).

#### **Collecting Fault Information**

**Step 13** On the FusionInsight Manager, choose **O&M** > **Log** > **Download**.

**Step 14** Select the following nodes in the required cluster from the **Service** drop-down list:




- KrbServer

 **NOTE**

KrbServer logs do not need to be downloaded in normal mode.

- ZooKeeper
- Storm

**Step 15** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 16** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.218 ALM-26052 Number of Available Supervisors of the Storm Service Is Less Than the Threshold

## Description

The system periodically checks the number of available Supervisors every 60 seconds and compares the number of available Supervisors with the threshold. This alarm is generated when the number of available Supervisors is less than the threshold.

You can change the threshold in **O&M > Alarm > Thresholds > Name of the desired cluster**.

This alarm is cleared when the number of available Supervisors is greater than or equal to the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
26052	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

Existing tasks in the cluster cannot be performed. The cluster can receive new Storm tasks, but cannot perform these tasks.

## Possible Causes

The status of some Supervisors in the cluster is abnormal.

## Procedure


### Check the Supervisor status.

- Step 1** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Supervisor** to go to the Storm service management page.
- Step 2** In **Roles**, check whether any instance whose status is **Faulty** or **Restoring** exists.
  - If yes, go to **Step 3**.
  - If no, go to **Step 5**.
- Step 3** Select Supervisor role instances whose status is **Faulty** or **Restoring**, choose **More** > **Restart Instance**, and check whether the instances restart successfully.
  - If yes, go to **Step 4**.
  - If no, go to **Step 5**.
- Step 4** Wait for 30 seconds, and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 5**.

 **NOTE**

Services are interrupted when the Supervisor is being restarted. Then, services are restored after the restarting.

**Collect fault information.**

- Step 5** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 6** Select **Storm** and **ZooKeeper** in the required cluster from the **Service** drop-down list box.
- Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 8** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.219 ALM-26053 Storm Slot Usage Exceeds the Threshold

### Description

The system checks the slot usage every 60 seconds and compares the actual slot usage with the threshold. This alarm is generated when the slot usage is greater than the threshold.

You can change the threshold in **O&M > Alarm > Thresholds**.

This alarm is cleared when the slot usage is less than or equal to the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
26053	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

New Storm tasks cannot be performed.

## Possible Causes

- The status of some Supervisors in the cluster is abnormal.
- The status of all Supervisors is normal, but the processing capability is insufficient.

## Procedure

### Check the Supervisor status.

- Step 1** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Instance** to go to the Storm instance management page.
- Step 2** Check whether any instance whose status is **Faulty** or **Restoring** exists.
  - If yes, go to **Step 3**.
  - If no, go to **Step 5**.
- Step 3** Select Supervisor role instances whose status is **Faulty** or **Restoring**, choose **More** > **Restart Instance**, and check whether the instances restart successfully.
  - If yes, go to **Step 4**.
  - If no, go to **Step 10**.
- Step 4** Wait several minutes, and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 5**.

### Increase the number of slots in each Supervisor.

- Step 5** Log in to the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Configurations** > **All Configurations**.
- Step 6** Increase the number of ports in the **supervisor.slots.ports** parameter of each Supervisor role and restart the instance.

**Step 7** Wait several minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Step 8** Perform capacity expansion for Supervisor.

**Step 9** Wait several minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).


 **NOTE**

Services are interrupted when the Supervisor is being restarted. Then, services are restored after the restarting.

**Collect fault information.**

**Step 10** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 11** Select **Storm** and **ZooKeeper** in the required cluster from the **Service** drop-down list box.

**Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.220 ALM-26054 Nimbus Heap Memory Usage Exceeds the Threshold

### Description

The system checks the heap memory usage of Storm Nimbus every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Storm Nimbus exceeds the threshold (80% of the maximum memory by default) for 5 consecutive times.

Users can choose **O&M > Alarm > Thresholds > Name of the desired cluster > Storm > Nimbus** to change the threshold.

The alarm is cleared when the heap memory usage is less than or equal to the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
26054	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

When the heap memory usage of Storm Nimbus is overhigh, frequent GCs occur. In addition, a memory overflow may occur so that the Yarn service is unavailable.

## Possible Causes

The heap memory of the Storm Nimbus instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

## Procedure

**Check the heap memory usage.**

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Heap Memory Usage of Storm Nimbus Exceeds the Threshold > Location**. Check the host name of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Storm > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize > Nimbus > Heap Memory Usage of Nimbus**. Click **OK**.

**Step 3** Check whether the used heap memory of Nimbus reaches the threshold (The default value is 80% of the maximum heap memory) specified for Nimbus.

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Configurations** > **All Configurations** > **Nimbus** > **System**. Change the value of **-Xmx** in **NIMBUS\_GC\_OPTS** based on site requirements, and click **Save**. Click **OK**.

 **NOTE**

- You are advised to set **-Xms** and **-Xmx** to the same value to prevent adverse impact on performance when JVM dynamically adjusts the heap memory size.
- The number of Workers grows as the Storm cluster scale increases. You can increase the value of **GC\_OPTS** for Nimbus. The recommended value is as follows: If the number of Workers is 20, set **-Xmx** to a value greater than or equal to 1 GB. If the number of Workers exceeds 100, set **-Xmx** to a value greater than or equal to 5 GB.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

 **NOTE**


During the restart of the service or instance, services are interrupted. After the service or instance is restarted, services are restored.

### Collect fault information.

**Step 6** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

**Step 7** Select the following node in the required cluster from the **Service** drop-down list.

- NodeAgent
- Storm

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.221 ALM-27001 DBService Service Unavailable

### Description

The alarm module checks the DBService service status every 30 seconds. This alarm is generated when the system detects that DBService service is unavailable.

This alarm is cleared when DBService service recovers.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
27001	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

### Impact on the System

The database service is unavailable and cannot provide data import and query functions for upper-layer services, which results in some services exceptions.

### Possible Causes

- The floating IP address does not exist.
- There is no active DBServer instance.
- The active and standby DBServer processes are abnormal.

### Procedure

**Check whether the floating IP address exists in the cluster environment.**

**Step 1** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService** > **Instance**.

**Step 2** Check whether the active instance exists.



- If yes, go to [Step 3](#).
- If no, go to [Step 9](#).

**Step 3** Select the active DBServer instance and record the IP address.

**Step 4** Log in to the host that corresponds to the preceding IP address as user **root**, and run the **ifconfig** command to check whether the DBService floating IP address exists on the node.

- If yes, go to [Step 5](#).
- If no, go to [Step 9](#).

**Step 5** Run the **ping floatip** command to check whether the DBService floating IP address can be pinged successfully.

- If yes, go to [Step 6](#).
- If no, go to [Step 9](#).

**Step 6** Log in to the host that corresponds to the DBService floating IP address as user **root**, and run the command to delete the floating IP address.

**ifconfig interface down**

**Step 7** On the FusionInsight Manager home page, choose **Cluster > Services > DBService > More > Restart Service** to restart DBService, and check whether DBService is restarted successfully.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

- If yes, go to [Step 8](#).
- If no, go to [Step 9](#).

**Step 8** Wait for about 2 minutes and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to [Step 14](#).

**Check the status of the active DBServer instance.**

**Step 9** Select the DBServer instance whose role status is abnormal and record the IP address.

**Step 10** On the **Alarm** page, check whether **Process Fault** occurs in the DBServer instance on the host that corresponds to the IP address.

- If yes, go to [Step 11](#).
- If no, go to [Step 14](#).

**Step 11** Handle the alarm according to "ALM-12007 Process Fault".

**Step 12** Wait for about 5 minutes and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to [Step 19](#).

**Check the status of the active and standby DBServers.**

- Step 13** Log in to the host that corresponds to the preceding IP address as user **root**, and run the **su - omm** command to switch to user **omm**.
- Step 14** Run the **cd \${DBSERVER\_HOME}** command to go to the installation directory of the DBService.
- Step 15** Run the **sh sbin/status-dbserver.sh** command to view the status of the active and standby HA processes of DBService. Determine whether the status can be viewed successfully.

```

HAMode
double

NodeName HostName HAVersion StartTime HAActive
HAAllResOK HARunPhase
10_5_89_12 host01 V100R001C01 2019-06-13 21:33:09 active
normal Activated
10_5_89_66 host03 V100R001C01 2019-06-13 21:33:09 standby
normal Deactivated

NodeName ResName ResStatus ResHASStatus ResType
10_5_89_12 floatip Normal Normal Single_active
10_5_89_12 gaussDB Active_normal Normal Active_standby
10_5_89_66 floatip Stopped Normal Single_active
10_5_89_66 gaussDB Standby_normal Normal Active_standby

```


- If yes, go to [Step 16](#).
  - If no, go to [Step 19](#).
- Step 16** Check whether the active and standby HA processes are in the abnormal state.
- If yes, go to [Step 17](#).
  - If no, go to [Step 19](#).
- Step 17** On FusionInsight Manager, choose **Cluster > Services > DBService > More > Restart Service** to restart DBService, and check whether the system displays a message indicating that the restart is successful.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

- If yes, go to [Step 18](#).
  - If no, go to [Step 19](#).
- Step 18** Wait for about 2 minutes and check whether the alarm is cleared in the alarm list.
- If yes, no further action is required.
  - If no, go to [Step 19](#).

**Collect fault information.**

- Step 19** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 20** Select **DBService** in the required cluster and **NodeAgent** from the **Service**.
- Step 21** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 22** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.222 ALM-27003 DBService Heartbeat Interruption Between the Active and Standby Nodes

### Description

This alarm is generated when the active or standby DBService node does not receive heartbeat messages from the peer node for 7 seconds.

This alarm is cleared when the heartbeat recovers.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
27003	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local DBService HA Name	Specifies a local DBService HA.
Peer DBService HA Name	Specifies a peer DBService HA.

### Impact on the System


During the DBService heartbeat interruption, only one node can provide the service. If this node is faulty, no standby node is available for failover and the service is unavailable.

## Possible Causes

The link between the active and standby DBService nodes is abnormal.

## Procedure

**Check whether the network between the active DBService server and the standby DBService server is normal.**

**Step 1** In the alarm list on FusionInsight Manager, click  in the row where the alarm is located in the real-time alarm list and view the standby DBService server address.

**Step 2** Log in to the active DBService server as user **root**.

**Step 3** Run the **ping *standby DBService heartbeat IP address*** command to check whether the standby DBService server is reachable.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

**Step 4** Contact the network administrator to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

**Step 5** Rectify the network fault and check whether the alarm is cleared from the alarm list.


- If yes, no further action is required.
- If no, go to **Step 6**.

**Collect fault information.**

**Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 7** Select the following nodes in the required cluster from the **Service**:

- DBService
- Controller
- NodeAgent

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.223 ALM-27004 Data Inconsistency Between Active and Standby DBServices

### Description

The system checks the data synchronization status between the active and standby DBService every 10 seconds. This alarm is generated when the synchronization status cannot be queried for six consecutive times or when the synchronization status is abnormal.

This alarm is cleared when the synchronization status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
27004	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local DBService HA Name	Specifies the HA name of the local DBService.
Peer DBService HA Name	Specifies the HA name of the peer DBService.
SYNC_PERCENT	Specifies the synchronization percentage.

### Impact on the System

When data is not synchronized between the active and standby DBServices, data may be lost or abnormal if the active instance becomes abnormal.

## Possible Causes

- The network between the active and standby nodes is unstable.
- The standby DBService is abnormal.
- The standby node disk space is full.
- The CPU usage of the GaussDB process on the active DBService node is high. You need to locate the failure cause based on logs.

## Procedure

### Check whether the network between the active and standby nodes is normal.

- Step 1** On FusionInsight Manager, choose **Cluster > Services > DBService > Instance**, check the service IP address of the standby DBServer instance.
- Step 2** Log in to the active DBService node as user **root**.
- Step 3** Run the **ping Standby DBService heartbeat IP address** command to check whether the standby DBService node is reachable.
- If yes, go to **Step 6**.
  - If no, go to **Step 4**.
- Step 4** Contact the network administrator to check whether the network is faulty.
- If yes, go to **Step 5**.
  - If no, go to **Step 6**.
- Step 5** Rectify the network fault and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 6**.

### Check whether the standby DBService is normal.

- Step 6** Log in to the standby DBService node as user **root**.
- Step 7** Run the **su - omm** command to switch to user **omm**.
- Step 8** Go to the **`\${DBSERVER\_HOME}/sbin** directory and run the **./status-dbserver.sh** command to check whether the GaussDB resource status of the standby DBService is normal. In the command output, check whether the following information is displayed in the row where **ResName** is **gaussDB**:

For example:

```
10_10_10_231 gaussDB Standby_normal Normal Active_standby
```

- If yes, go to **Step 9**.
- If no, go to **Step 16**.

**Check whether the standby node disk space is full.** (Skip this check for versions later than MRS 3.1.2.)

- Step 9** Log in to the standby DBService node as user **root**.
- Step 10** Run the **su - omm** command to switch to user **omm**.
- Step 11** Go to the **`\${DBSERVER\_HOME}** directory, and run the following commands to obtain the DBService data directory:

```
cd ${DBSERVER_HOME}
source .dbservice_profile
echo ${DBSERVICE_DATA_DIR}
```

**Step 12** Run the `df -h` command to view the system disk partition usage information.

**Step 13** Check whether the DBService data directory space is full.

- If yes, go to [Step 14](#).
- If no, go to [Step 16](#).

**Step 14** Expand the disk capacity.


**Step 15** After the disk capacity is expanded, wait 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 16](#).

**Collect fault information.**

**Step 16** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 17** In the **Service** area, select **DBService** of the target cluster and **OS, OS Statistics**, and **OS Performance** under **O&M**, and click **OK**.

**Step 18** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 19** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.224 ALM-27005 Database Connections Usage Exceeds the Threshold

### Description

The system checks the usage of the number of database connections of the nodes where DBServer instances are located every 30 seconds and compares the usage with the threshold. If the usage exceeds the threshold for five consecutive times (this number is configurable, and 5 is the default value), the system generates this alarm. The default usage threshold is 90%, and you can configure it based on site requirements.

The trigger count is configurable. This alarm is cleared in the following scenarios:

- The trigger count is 1, and the usage of the number of database connections is less than or equal to the threshold.
- The trigger count is greater than 1, and the usage of the number of database connections is less than or equal to 90% of the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
27005	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

Upper-layer services may fail to connect to the DBService database, affecting services.

## Possible Causes

- Too many database connections are used.
- The maximum number of database connections is improperly configured.
- The alarm threshold or alarm trigger count is improperly configured.

## Procedure


### Checking whether too many data connections are used

- Step 1** On FusionInsight Manager, click DBService in the service list on the left navigation pane. The DBService monitoring page is displayed.



**Step 2** Observe the number of connections used by the database user, as shown in [Figure 7-125](#). Based on the service scenario, reduce the number of database user connections.

**Figure 7-125** Number of connections used by database users




**Step 3** Wait for 2 minutes and check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to [Step 4](#).

### Checking whether the maximum number of database connections is properly configured

**Step 4** Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService** > **Configurations**. On the displayed page, select the **All Configurations** tab, and increase the maximum number of database connections based on service requirements, as shown in [Figure 7-126](#). Click **Save**. In the displayed **Save configuration** dialog box, click **OK**.

**Figure 7-126** Setting the maximum number of database connections



**Step 5** After the maximum number of database connections is changed, restart DBService (do not restart the upper-layer services).

**Procedure:** Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService**. On the displayed page, choose **More** > **Restart Service**. Enter the password of the current login user and click **OK**. Do not select **Restart upper-layer services**., click **OK**.

 **NOTE**

The service will be unavailable during the restart. In addition, upper-layer services that depend on the service are affected.

**Step 6** After the service is restarted, wait for 2 minutes and check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 7](#).

**Checking whether the alarm threshold or trigger count is properly configured**

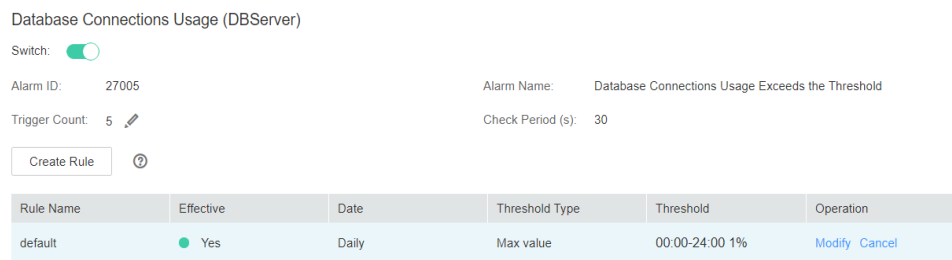
**Step 7** Log in to FusionInsight Manager and change the alarm threshold and alarm trigger count based on the actual database connection usage.

Choose **O&M > Alarm > Thresholds > Name of the desired cluster > DBService > Database > Database Connections Usage (DBServer)**. In the **Database Connections Usage (DBServer)** area, click the pencil icon next to **Trigger Count**. In the displayed dialog box, change the trigger count, as shown in [Figure 7-127](#).

 **NOTE**

**Trigger Count:** If the usage of the number of database connections exceeds the threshold consecutively for more than the value of this parameter, an alarm is generated.

**Figure 7-127** Setting alarm trigger count



Based on the actual database connection usage, choose **O&M > Alarm > Thresholds > Name of the desired cluster > DBService > Database > Database Connections Usage (DBServer)**. In the **Database Connections Usage (DBServer)** area, click **Modify** in the **Operation** column. In the **Modify Rule** dialog box, modify the required parameters and click **OK** as shown in [Figure 7-128](#).

**Figure 7-128** Set alarm thresholdThresholds > **Modify Rule**

---

\* Rule Name:

\* Severity:

\* Threshold Type:  Max value  Min value

\* Date:  Daily  
 Weekly  
 Other

Thresholds:	Start and End Time	Threshold
	<input type="text" value="00:00"/> - <input type="text" value="23:59"/>	<input type="text" value="90"/>

**Step 8** Wait for 2 minutes and check whether the alarm is automatically cleared.


- If it is, no further action is required.
- If it is not, go to **Step 9**.

#### Collect fault information

**Step 9** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 10** Select **DBService** in the required cluster from the **Service**.

**Step 11** Specify the host for collecting logs by setting the **Host** parameter that is optional. By default, all hosts are selected.

**Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact the O&M personnel and send the collected fault logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.225 ALM-27006 Disk Space Usage of the Data Directory Exceeds the Threshold

### Description

The system checks the disk space usage of the data directory on the active DBServer node every 30 seconds and compares the disk usage with the threshold. The alarm is generated when the disk space usage exceeds the threshold for five consecutive times (the default value). The number of consecutive times is configurable. The disk space usage threshold of the data directory is set to 80% by default, which is configurable as well.

The value of **hit number** is configurable. When the value is set to **1** and the disk space usage is lower than or equal to the threshold, the alarm is cleared. When the value is greater than 1 and the disk space usage is lower than 90% of the threshold, the alarm is cleared.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
27006	Major	Yes

### Parameters

Name	Meaning
ClusterName	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition where the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the actual indicator value exceeds this threshold, the alarm is generated.

### Impact on the System

- The DBService service process cannot provide the API for data writing.

- When the disk space usage of the data directory exceeds 90%, the database enters the read-only mode and "Database Enters the Read-Only Mode" is generated. As a result, service data cannot be written to the database.

## Possible Causes

- The alarm threshold is improperly configured.
- The data volume of the database is too large or the disk configuration cannot meet service requirements, causing excessive disk usage.

## Procedure

### Check whether the threshold is set properly.

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Name of the desired cluster > DBService > Database > Disk Space Usage of the Data Directory** to check whether the alarm threshold is proper (the default value 80% is a proper value).


- If yes, go to [Step 3](#).
- If no, go to [Step 2](#).

**Step 2** Change the alarm threshold based on the actual service situation.

**Step 3** Choose **Cluster > Name of the desired cluster > Services > DBService**. On the **Dashboard** page, view the **Disk Space Usage of the Data Directory** chart and check whether the disk space usage of the data directory is lower than the threshold.

- If yes, go to [Step 4](#).
- If no, go to [Step 5](#).

**Figure 7-129** Disk Space Usage of the Data Directory



**Step 4** Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

### Check whether large files are incorrectly written into the disk.

**Step 5** Log in to the active DBService node as user **omm**.

**Step 6** Run the following commands to view the files whose size exceeds 500 MB in the data directory and check whether there are large files incorrectly written into the directory:

```
source $DBSERVER_HOME/.dbservice_profile
```

```
find "$DBSERVICE_DATA_DIR"/../ -type f -size +500M
```

- If yes, go to [Step 7](#).
- If no, go to [Step 8](#).

**Step 7** Handle the large files based on the actual scenario and check whether the alarm is cleared 2 minutes later.


- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect fault information.**

**Step 8** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 9** Expand the **Service** drop-down list, and select **DBService** for the target cluster.

**Step 10** Specify the host for collecting logs by setting the **Host** parameter which is optional. By default, all hosts are selected.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.226 ALM-27007 Database Enters the Read-Only Mode

### Description

The system checks the disk space usage of the data directory on the active DBServer node every 30 seconds. The alarm is generated when the disk space usage exceeds 90%.

The alarm is cleared when the disk space usage is lower than 80%.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
27007	Critical	Yes

## Parameters

Name	Meaning
ClusterName	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the actual indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

- Service data is lost.
- Data cannot be written for upper-layer services and the data is lost.

## Possible Causes

The disk configuration cannot meet service requirements. The disk usage reaches the upper limit.

## Procedure

**Check whether the disk space usage reaches the upper limit.**

- Step 1** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService**.
- Step 2** On the **Dashboard** page, view the **Disk Space Usage of the Data Directory** chart and check whether the disk space usage of the data directory exceeds 90%.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 13](#).
- Step 3** Log in to the active management node of the DBServer as user **omm** and run the following commands to check whether the database enters the read-only mode:
 

```
source $DBSERVER_HOME/.dbservice_profile
gsq -U omm -W password -d postgres -p 20051
```

```
show default_transaction_read_only;
```

**NOTE**

In the preceding commands, *password* indicates the password of user **omm** of the DBService database (You can view the initial password of user omm in User [User Account List](#)). You can run the `\q` command to exit the database.

Check whether the value of **default\_transaction\_read\_only** is **on**.

```
POSTGRES=# show default_transaction_read_only;
default_transaction_read_only

on
(1 row)
```

- If yes, go to [Step 4](#).
- If no, go to [Step 13](#).

**Step 4** Run the following commands to open the **dbservice.properties** file:

```
source $DBSERVER_HOME/.dbservice_profile
```

```
vi ${DBSERVICE_SOFTWARE_DIR}/tools/dbservice.properties
```

**Step 5** Change the value of **gaussdb\_readonly\_auto** to **OFF**.

**Step 6** Run the following command to open the **postgresql.conf** file:

```
vi ${DBSERVICE_DATA_DIR}/postgresql.conf
```

**Step 7** Delete **default\_transaction\_read\_only = on**.

**Step 8** Run the following command for the configuration to take effect:

```
gs_ctl reload -D ${DBSERVICE_DATA_DIR}
```

**Step 9** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. On the right of the alarm "Database Enters the Read-Only Mode", click **Clear** in the **Operation** column. In the dialog box that is displayed, click **OK** to manually clear the alarm.

**Step 10** Log in to the active management node of the DBServer as user **omm** and run the following commands to view the files whose size exceeds 500 MB in the data directory and check whether there are large files incorrectly written into the directory:

```
source $DBSERVER_HOME/.dbservice_profile
```

```
find "$DBSERVICE_DATA_DIR"/../ -type f -size +500M
```

- If yes, go to [Step 11](#).
- If no, go to [Step 13](#).

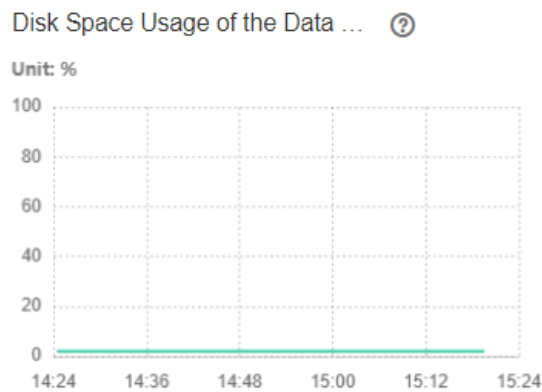
**Step 11** Handle the files that are incorrectly written into the directory based on the actual scenario.

**Step 12** Log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > DBService**. On the **Dashboard** page, view the **Disk Space Usage of the Data Directory** chart and check whether the disk space usage is lower than 80%.




- If yes, no further action is required.
- If no, go to [Step 13](#).

**Figure 7-130** Disk Space Usage of the Data Directory



**Collect fault information.**

- Step 13** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 14** Expand the **Service** drop-down list, and select **DBService** for the target cluster.
- Step 15** Specify the host for collecting logs by setting the **Host** parameter which is optional. By default, all hosts are selected.
- Step 16** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 17** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.227 ALM-29000 Impala Service Unavailable

### Alarm Description

The alarm module checks the Impala service status every 30 seconds. This alarm is generated if the Impala service is abnormal.

This alarm is cleared after the Impala service recovers.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29000	Critical	Yes

## Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

## Impact on the System

When the Impala service is abnormal, you cannot perform cluster operations on Impala through FusionInsight Manager. The Impala service functions are unavailable.

## Possible Causes

- The Hive service is abnormal.
- The KrbServer service is abnormal.
- The Impala process is abnormal.

## Handling Procedure

**Check whether the services on which Impala depends are normal.**

**Step 1** On FusionInsight Manager, choose **Cluster > Services** to check whether Hive and KrbServer are stopped.

- If yes, start the stopped services and go to [Step 2](#).
- If no, go to [Step 3](#).

**Step 2** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the alarm list, check whether the **Impala Service Unavailable** alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

- Step 3** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the alarm list, check whether ALM-16004 Hive Service Unavailable and ALM-25500 KrbServer Service Unavailable exist.
- If yes, go to **Step 4**.
  - If no, go to **Step 5**.
- Step 4** Rectify the fault by following the handling procedure of **ALM-16004 Hive Service Unavailable** or **ALM-25500 KrbServer Service Unavailable**. Then, check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 5**.
- Check whether the Impala process is normal.**
- Step 5** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Check whether ALM-12007 Process Fault exists in the alarm list.
- If yes, go to **Step 6**.
  - If no, go to **Step 7**.
- Step 6** Rectify the fault by referring to the handling method of **ALM-12007 Process Fault**, and then check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 7**.
- Collect the fault information.**
- Step 7** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 8** Expand the **Service** drop-down list, and select **Impala** for the target cluster.
- Step 9** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

## Related Information

None

## 7.12.228 ALM-29004 Impalad Process Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the memory usage of the Impalad process every 30 seconds. This alarm is generated when the system detects that the memory usage exceeds the default threshold (80%).

This alarm is automatically cleared when the system detects that the memory usage of the process falls below the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29004	Minor	Yes

### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System


The memory usage is too high. Some query tasks may fail due to insufficient memory.

### Possible Causes

The Impalad process is executing a large number of query tasks.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > CPU and Memory > Impalad Process Memory Usage (Impalad)** and check the threshold.



**Step 2** If the alarm threshold is smaller than 80%, increase the alarm threshold as required and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

**Step 3** If the threshold is greater than 80%, check whether a large number of concurrent query tasks exist when the alarm is generated. A large number of concurrent query tasks will cause the memory usage to increase sharply. After the tasks are complete, check whether the alarm is automatically cleared. During this period, some tasks may fail to be executed or may be canceled due to insufficient memory. In this case, try again.

### NOTE

If the memory usage always exceeds the threshold, the cluster capacity needs to be expanded.

- If yes, no further action is required.
- If no, go to [Step 4](#).

### Collect fault information.

**Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 5** Expand the **Service** drop-down list, and select **Impala** for the target cluster.

**Step 6** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 7** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

The alarm is automatically cleared after the burst concurrent tasks are complete.

## Related Information

None

## 7.12.229 ALM-29005 Number of JDBC Connections to Impalad Exceeds the Threshold

### Alarm Description

The system checks the number of client connections to the Impalad node every 30 seconds. This alarm is generated when the number of client connections exceeds the customized threshold (60 by default).

This alarm is automatically cleared when the number of client connections is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29005	Major	Yes

### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

New client connections may be blocked or even fail.

### Possible Causes

The number of client connections maintained by the Impalad service is too large or the threshold is too small.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > Connections > Number of JDBC Connections to Impalad Process** to check the configured threshold.

Number of JDBC Connections to Impalad Process (Impalad)

Switch:

Alarm ID: 29005      Alarm Name: Number of JDBC Connections to Impalad Exceeds the Threshold

Trigger Count: 1      Check Period (s): 30

[Create Rule](#)

Rule Name	Effective	Date	Threshold Type	Threshold	Operation
default	<span style="color: green;">●</span> Yes	Daily	Max value	00:00-24:00 <b>60</b>	<a href="#">Modify</a> <a href="#">Cancel</a>

**Step 2** Check the number of JDBC applications connected to Impalad and stop idle applications. Check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 3** to change the number of concurrent client connections.

**Step 3** On FusionInsight Manager, choose **Cluster > Impala > Configurations > All Configurations > Impalad > Customization**. Add the customized parameter -- **fe\_service\_threads**. The default value of this parameter is **64**. Change the value as required and click **Save**.

[Save](#) [Import](#) [Export](#)

Basic Configurations [All Configurations](#)

- Impala
  - Customization
  - HDFSClient
  - OBS
  - Ranger
- StateStore
- Catalog
- Impalad
  - Customization**
  - Environment
  - Log Configuration
  - Ranger
  - System

Parameter	Value				
impalad.customized.configs	<table border="1"> <thead> <tr> <th>Name</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>--fe_service_threads</td> <td><input type="text" value="60"/></td> </tr> </tbody> </table>	Name	Value	--fe_service_threads	<input type="text" value="60"/>
Name	Value				
--fe_service_threads	<input type="text" value="60"/>				

**Step 4** After the query tasks on all clients are complete, click the **Instance** tab. Select all Impalad instances, and restart them.

### NOTE

The service will become unavailable when all instances are restarted. If a single instance is restarted, the tasks that are being executed on that instance will fail and the service will become available.

Dashboard [Chart](#) [Instance](#) [Instance Groups](#) [Configurations](#)

[Add Instance](#) [Start Instance](#) [More](#)

Role	Configuration Status	Host Name	Management IP Address	Service IP Address
<input type="checkbox"/> Catalog	<span style="color: green;">●</span> Synchronized	node-master1fJZ	192.168.0.51	192.168.0.51
<input checked="" type="checkbox"/> Impalad	<span style="color: green;">●</span> Synchronized	node-group-1WDQ...	192.168.0.113	192.168.0.113
<input checked="" type="checkbox"/> Impalad	<span style="color: green;">●</span> Synchronized	node-group-1WDQ...	192.168.0.208	192.168.0.208
<input checked="" type="checkbox"/> Impalad	<span style="color: green;">●</span> Synchronized	node-group-1WDQ...	192.168.0.231	192.168.0.231
<input type="checkbox"/> StateStore	<span style="color: green;">●</span> Synchronized	node-master1fJZ	192.168.0.51	192.168.0.51

**Step 5** After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Impala** for the target cluster.

**Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

## Related Information

None

## 7.12.230 ALM-29006 Number of ODBC Connections to Impalad Exceeds the Threshold

### Alarm Description

The system checks the number of client connections to the Impalad node every 30 seconds. This alarm is generated when the number of client connections exceeds the customized threshold (60 by default).

This alarm is automatically cleared when the number of client connections is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29006	Major	Yes



## Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System


New client connections may be blocked or even fail.

## Possible Causes

The number of client connections maintained by the Impalad service is too large or the threshold is too small.

## Handling Procedure


- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > Connections > Number of ODBC Connections to Impalad Process (Impalad)** to check the threshold.



- Step 2** Check the number of ODBC applications connected to Impalad and stop idle applications. Check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 3** to change the number of concurrent connections supported by Impalad.

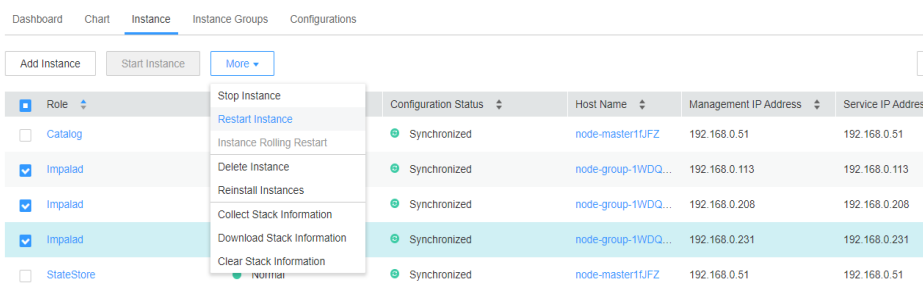
- Step 3** On FusionInsight Manager, choose **Cluster > Impala > Configurations > All Configurations > Impalad > Customization**. Add the customized parameter -- **fe\_service\_threads**. The default value of this parameter is **64**. Change the value as required and click **Save**.



**Step 4** After the query tasks on all clients are complete, click the **Instance** tab. Select all Impalad instances, and restart them.

**NOTE**

The service will become unavailable when all instances are restarted. If a single instance is restarted, the tasks that are being executed on that instance will fail and the service will become available.



**Step 5** After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Impala** for the target cluster.

**Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

**Alarm Clearance**

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

## Related Information

None

### 7.12.231 ALM-29007 Impalad Process Memory Usage Exceeds the Threshold

#### Alarm Description

The system checks the memory usage of the Impalad process every 30 seconds. This alarm is generated when the system detects that the memory usage exceeds the default threshold (80%).

This alarm is automatically cleared when the system detects that the memory usage of the process falls below the threshold.

#### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29007	Minor	Yes

#### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

#### Impact on the System

The memory usage is too high. Some query tasks may fail due to insufficient memory.

#### Possible Causes

The Impalad process is executing a large number of query tasks.

## Handling Procedure

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > CPU and Memory > Impalad Process Memory Usage (Impalad)** and check the threshold.
- Step 2** If the alarm threshold is smaller than 80%, increase the alarm threshold as required and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 3](#).
- Step 3** If the threshold is greater than 80%, check whether a large number of concurrent query tasks exist when the alarm is generated. A large number of concurrent query tasks will cause the memory usage to increase sharply. After the tasks are complete, check whether the alarm is automatically cleared. During this period, some tasks may fail to be executed or may be canceled due to insufficient memory. In this case, try again.

### NOTE

If the memory usage always exceeds the threshold, the cluster capacity needs to be expanded.

- If yes, no further action is required.
- If no, go to [Step 4](#).

### Collect fault information.

- Step 4** On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 5** Expand the **Service** drop-down list, and select **Impala** for the target cluster.
- Step 6** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 7** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

The alarm is automatically cleared after the burst concurrent tasks are complete.

## Related Information

None

## 7.12.232 ALM-29008 Number of ODBC Connections to Impalad Exceeds the Threshold

### Alarm Description

The system checks the number of client connections to the Impalad node every 30 seconds. This alarm is generated when the number of client connections exceeds the customized threshold (60 by default).

This alarm is automatically cleared when the number of client connections is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29008	Major	Yes

### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

New client connections may be blocked or even fail.

### Possible Causes

The number of client connections maintained by the Impalad service is too large or the threshold is too small.

## Handling Procedure

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > Connections > Number of ODBC Connections to Impalad Process (Impalad)** to check the threshold.
- Step 2** Check the number of ODBC applications connected to Impalad and stop idle applications. Check whether the alarm is automatically cleared.
- If yes, no further action is required.
  - If no, go to **Step 3** to change the number of concurrent connections supported by Impalad.
- Step 3** On FusionInsight Manager, choose **Cluster > Impala > Configurations > All Configurations > Impalad > Customization**. Add the custom parameter **--fe\_service\_threads**. The default value of this parameter is **64**. Change the value as required and click **Save**.
- Step 4** After the query tasks on all clients are complete, click the **Instances** tab. Select all Impalad instances, and restart them.

### NOTE

The service will become unavailable when all instances are restarted. If a single instance is restarted, the tasks that are being executed on that instance will fail and the service will become available.

- Step 5** After the restart is complete, check whether the alarm is cleared.
- If yes, no further action is required.
  - If yes, go to **Step 6**.

### Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Impala** for the target cluster.
- Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.233 ALM-29010 Number of Queries Being Submitted by Impalad Exceeds the Threshold

### Alarm Description

The system checks the total number of queries being submitted by the Impalad node every 60 seconds. This alarm is generated when the number of queries exceeds the customized threshold (150 by default).

This alarm is automatically cleared when the number of queries is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29010	Major	Yes

### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System


The queries may be blocked or even fail.

### Possible Causes


The Impalad service has maintained a large number of queries, or the threshold is too small.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > Query Task Sum Statistics > Total number of Queries Being Submitted (Impalad)** and check the threshold.




**Step 2** Change the threshold.



**Step 3** Click the **Instances** tab, select all Impalad instances, and restart them.

### NOTE

The service will become unavailable when all instances are restarted. If a single instance is restarted, the tasks that are being executed on that instance will fail and the service will become available.



**Step 4** After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

### Collect fault information.

**Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** Expand the **Service** drop-down list, and select **Impala** for the target cluster.



**Step 7** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.234 ALM-29011 Number of Queries Being Executed by Impalad Exceeds the Threshold

## Alarm Description

The system checks the total number of queries being executed by the Impalad node every 60 seconds. This alarm is generated when the number of queries exceeds the customized threshold (150 by default).

This alarm is automatically cleared when the number of queries is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29011	Major	Yes

## Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

The queries may be blocked or even fail.

## Possible Causes

The Impalad service has maintained a large number of queries, or the threshold is too small.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > Query Task Sum Statistics > Total number of Queries Being Executed (Impalad)** and check the threshold.

Total number of Queries Being Executed (Impalad)

Switch:

Alarm ID: 29011      Alarm Name: Total number of queries being executed to Impalad Exceeds the Threshold

Trigger Count: 1      Check Period (s): 60

Create Rule

Rule Name	Effective	Date	Threshold Type	Threshold	Operation
default	<span style="color: green;">●</span> Yes	Daily	Max value	00:00-24:00 <b>150</b>	<a href="#">Modify</a> <a href="#">Cancel</a>

**Step 2** Change the threshold.

Thresholds > **Modify Rule**

Rule Name: default

Alarm Severity: Major

Threshold Type:  Max value  Min value

Date:  Daily  Weekly  Others

Thresholds: Start and End Time      Threshold

00:00 - 23:59      **1**

**Step 3** Click the **Instances** tab, select all Impalad instances, and restart them.

### NOTE

The service will become unavailable when all instances are restarted. If a single instance is restarted, the tasks that are being executed on that instance will fail and the service will become available.

Dashboard   Chart   **Instance**   Instance Groups   Configurations

Add Instance   Start Instance   More

Role	Configuration Status	Host Name	Management IP Address	Service IP Address	Rack	Instance Group
Impalad	<span style="color: green;">●</span> Synchronized	node-madart0R1...	192.168.0.44	192.168.0.44	iddefaultrack1355	Catalog-DEFAULT
Impalad	<span style="color: green;">●</span> Synchronized	node-ana-corsKCF...	192.168.0.187	192.168.0.187	iddefaulttrackd	Impalad-DEFAULT
Impalad	<span style="color: green;">●</span> Synchronized	node-ana-corsKCF...	192.168.0.233	192.168.0.233	iddefaulttrackd	Impalad-DEFAULT
Impalad	<span style="color: green;">●</span> Synchronized	node-ana-corsKCF...	192.168.0.65	192.168.0.65	iddefaulttrackd	Impalad-DEFAULT
StateStore	<span style="color: green;">●</span> Synchronized	node-madart0R1...	192.168.0.44	192.168.0.44	iddefaultrack1355	StateStore-DEFAULT

**Step 4** After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

**Collect fault information.**

**Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** Expand the **Service** drop-down list, and select **Impala** for the target cluster.

**Step 7** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

**Alarm Clearance**

This alarm is automatically cleared after the fault is rectified.

**Related Information**

None

**7.12.235 ALM-29012 Number of Queries Being Waited by Impalad Exceeds the Threshold****Alarm Description**

The system checks the total number of queries being waited by the Impalad node every 60 seconds. This alarm is generated when the number of queries exceeds the customized threshold (150 by default).

This alarm is automatically cleared when the number of queries is less than the threshold.

**Alarm Attributes**

Alarm ID	Alarm Severity	Auto Cleared
29012	Major	Yes

**Alarm Parameters**

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.

Type	Parameter	Description
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System


The queries may be blocked or even fail.

## Possible Causes

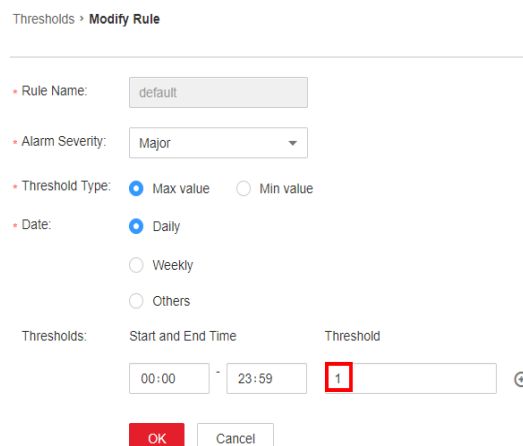
The Impalad service has maintained a large number of queries, or the threshold is too small.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > Query Task Sum Statistics > Total number of Waiting Queries (Impalad)** and check the threshold.




**Step 2** Change the threshold.



**Step 3** Click the **Instances** tab, select all Impalad instances, and restart them.

**NOTE**

The service will become unavailable when all instances are restarted. If a single instance is restarted, the tasks that are being executed on that instance will fail and the service will become available.



**Step 4** After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Collect fault information.**

**Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** Expand the **Service** drop-down list, and select **Impala** for the target cluster.

**Step 7** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

**Alarm Clearance**

This alarm is automatically cleared after the fault is rectified.

**Related Information**

None

**7.12.236 ALM-29013 Impalad FGC Time Exceeds the Threshold**

**Alarm Description**

The system checks the FGC time of the Impalad service every 60 seconds. This alarm is generated when the FGC time exceeds the threshold (12 seconds) for five consecutive times. This alarm is cleared when the FGC time is less than or equal to the threshold.

**Alarm Attributes**

Alarm ID	Alarm Severity	Auto Cleared
29013	Major	Yes

## Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Data read and write are affected.


## Possible Causes

The memory of the node instance is overused or the heap memory is inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

Check the GC time.

**Step 1** Choose **O&M > Alarm > Thresholds > Impala > Process FGCT > Process FGCT of Impalad (Impalad)**, and check the threshold (12s by default).



**Step 2** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and check whether the alarm whose **Alarm ID** is **29013** exists in the alarm list.

- If yes, go to **3**.
- If no, no further action is required.

**Step 3** On FusionInsight Manager, choose **Cluster > Impala**, click the **Instances** tab, select the Impalad instance for which the alarm is generated, then click the **Chart** tab, locate the **Process FGCT** chart, and check whether the FGC time is greater than the threshold in **1**.

- If yes, go to **4**.
- If no, go to **Step 5**.

**Step 4** Choose **O&M > Alarm > Thresholds > Impala > Process FGCT > Process FGCT of Impalad (Impalad)**, and change the threshold to a value less than the time obtained in **3**. Then, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

**Collect fault information.**

**Step 5** On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** Expand the **Service** drop-down list, and select **Impala** for the target cluster.

**Step 7** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.237 ALM-29014 Catalog FGC Time Exceeds the Threshold

### Alarm Description

The system checks the FGC time of the Catalog service every 60 seconds. This alarm is generated when the FGC time exceeds the threshold (12 seconds) for five consecutive times. This alarm is cleared when the FGC time is less than or equal to the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29014	Major	Yes

### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.

Type	Parameter	Description
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Data read and write are affected.


## Possible Causes

The memory of the node instance is overused or the heap memory is inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

Check the GC time.

- Step 1** Choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, choose **Impala > Process FGCT > Process FGCT of Catalog (Catalog)**, and check the threshold (12s by default).



- Step 2** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and check whether the alarm whose **Alarm ID** is **29014** exists in the alarm list.

- If yes, go to **3**.
- If no, no further action is required.

- Step 3** On FusionInsight Manager, choose **Cluster > Impala**, click the **Instance** tab, select the Catalog instance for which the alarm is generated, then click the **Chart** tab, locate the **Process FGCT** chart, and check whether the FGC time is greater than the threshold in **1**.

- If yes, go to **4**.
- If no, go to **5**.

- Step 4** Choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, choose **Impala > Process FGCT > Process FGCT of Catalog (Catalog)**, and change the threshold to a value less than the time obtained in **3**. Then, check whether the alarm is cleared.



- If yes, no further action is required.
- If no, go to 5.

**Collect fault information.**

**Step 5** On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** Expand the **Service** drop-down list, and select **Impala** for the target cluster.

**Step 7** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.238 ALM-29015 Catalog Process Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the memory usage of the Catalog process every 30 seconds. This alarm is generated when the system detects that the memory usage exceeds the default threshold (80%).

This alarm is automatically cleared when the system detects that the memory usage of the process falls below the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29015	Major	Yes

### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.

Type	Parameter	Description
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

The memory usage is too high. Some query tasks may fail due to insufficient memory.

## Possible Causes

The memory of the node instance is overused or the memory is inappropriately configured.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Thresholds > Impala > CPU and Memory > Catalog Process Memory Usage (Impalad)** and check the threshold.

**Step 2** If the alarm threshold is smaller than 80%, increase the alarm threshold as required and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

**Step 3** If the threshold is greater than 80%, check whether a large number of concurrent query tasks exist when the alarm is generated. A large number of concurrent query tasks will cause the memory usage to increase sharply. After the tasks are complete, check whether the alarm is automatically cleared. During this period, some tasks may fail to be executed or may be canceled due to insufficient memory. In this case, try again.

### NOTE

If the memory usage always exceeds the threshold, the cluster capacity needs to be expanded.

- If yes, no further action is required.
- If no, go to [Step 4](#).

### Collect fault information.

**Step 4** On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 5** Expand the **Service** drop-down list, and select **Impala** for the target cluster.
- Step 6** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 7** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.239 ALM-29016 Impalad Instance in the Sub-healthy State

### Alarm Description

In MRS 3.1.5, the system checks every 60 seconds whether the Hive Server2 HTTP port (28000) of Impalad responds to cURL requests. This alarm is generated when the returned result has been incorrect for 20 seconds in two consecutive times. This alarm is cleared when the system correctly responds within 20 seconds.

In other MRS versions, the system checks every 60 seconds whether Impalad can execute **select 1**. This alarm is generated when the returned result has been incorrect for 20 seconds in two consecutive times. This alarm is cleared when the SQL statement is correctly executed within 20 seconds.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29016	Minor	Yes

### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.

Type	Parameter	Description
	HostName	Specifies the host for which the alarm was generated.

## Impact on the System

Impalad cannot execute SQL statements or SQL statement execution times out, which affects data read and write.

## Possible Causes

The Impalad service maintains too many queries.


## Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **Cluster > Services > Impala > Impalad Web UI**. On the displayed page, click any node to go to the web UI.
- Step 2** On the web UI, click **/backends** to view the Impala instance list. Locate the instance for which the alarm is generated and click **Web UI**. After the web UI of the subhealthy node is displayed, click **/queries** to check the task execution status and check whether any task is executed slowly.
  - If yes, go to **Step 3**.
  - If no, go to **Step 4**.
- Step 3** After the task is complete, check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 4**.
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Impala > Instances**, select the Impala instance for which the alarm is generated, click **More**, and select **Restart Instance**. Then, check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 5**.

### NOTE

The service will become unavailable when all instances are restarted. If a single instance is restarted, the tasks that are being executed on that instance will fail and the service will become available.

### Collect fault information.

- Step 5** On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 6** Expand the **Service** drop-down list, and select **Impala** for the target cluster.
- Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.240 ALM-29100 Kudu Service Unavailable

### Alarm Description

The system checks the Kudu service status every 60 seconds. This alarm is generated when the system detects that all Kudu instances are abnormal and considers that the Kudu service is unavailable.

This alarm is cleared when at least one Kudu instance becomes normal and the system considers that the Kudu instance service is restored.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29100	Critical	Yes

### Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

### Impact on the System

Users cannot use the Kudu service.

## Possible Causes


Some Kudu instances are abnormal.

## Handling Procedure

**Handle the Kudu instance exceptions.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the displayed page, locate the alarm ALM-29100 Kudu Service Unavailable.
- Step 2** In the **Location Information** column, record the host name and role name.
- Step 3** Choose **Cluster > Services > Kudu > Instances**. Click the role name for the host name obtained in **Step 2**, view the instance logs, and restore the instance. Then, check whether the alarm is cleared.
- If yes, go to **Step 4**.
  - If no, go to **Step 5**.
- Step 4** Choose **O&M > Alarm > Alarms** and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 5**.

**Collect the fault information.**

- Step 5** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 6** Expand the **Service** drop-down list, and select **Kudu** for the target cluster.
- Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

## Related Information

None

## 7.12.241 ALM-29104 Tserver Process Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the memory usage of the Kudu Tserver process every 60 seconds. This alarm is generated when the system detects that the memory usage of the Kudu Tserver process exceeds the threshold.

This alarm is cleared when the memory usage of the Tserver process becomes normal and the system considers that the Kudu instance service recovers.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29104	Critical	Yes

## Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Users cannot use the Kudu service.

## Possible Causes


The memory usage of a KuduTserver instance is too high.

## Handling Procedure

**Handle the Kudu instance exceptions.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the displayed page, locate alarm **ALM-29104 Tserver Process Memory Usage Exceeds the Threshold** and view the alarm source.
- Step 2** Choose **O&M > Alarm > Thresholds > Kudu**. Locate the threshold of this alarm and check whether the Kudu instance memory usage exceeds the threshold. If yes, rectify the fault or change the threshold.
- Step 3** Choose **O&M > Alarm** and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [4](#).

**Collect fault information.**

- Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 5** Expand the **Service** drop-down list, and select **Kudu** for the target cluster.
- Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 7** Contact O&M personnel and provide the collected logs.

----End

**Alarm Clearance**

This alarm is automatically cleared after the fault is rectified.

**Related Information**

None

## 7.12.242 ALM-29106 Tserver Process CPU Usage Exceeds the Threshold

**Alarm Description**

The system checks the Kudu service status every 60 seconds. This alarm is generated when the system detects that the CPU usage of the Kudu Tserver process is too high.

This alarm is cleared when the CPU usage of the Tserver process becomes normal and the system considers that the Kudu instance service recovers.

**Alarm Attributes**

Alarm ID	Alarm Severity	Auto Cleared
29106	Critical	Yes

**Alarm Parameters**

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.



Type	Parameter	Description
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Users cannot use the Kudu service.

## Possible Causes


The CPU usage of a KuduTserver instance is too high.

## Handling Procedure

**Handle the Kudu instance exceptions.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the displayed page, locate alarm **ALM-29106 Tserver Process CPU Usage Exceeds the Threshold** and view the alarm source.
- Step 2** Choose **O&M > Alarm > Thresholds > Kudu**. Locate the alarm threshold and check whether the CPU usage of the cluster Kudu instance exceeds the threshold. If yes, rectify the fault or change the threshold.
- Step 3** Choose **O&M > Alarm** and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [4](#).

**Collect fault information.**

- Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 5** Expand the **Service** drop-down list, and select **Kudu** for the target cluster.
- Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 7** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.243 ALM-29107 Tserver Process Memory Usage Exceeds the Threshold

## Alarm Description

The system checks the Kudu service status every 60 seconds. This alarm is generated when the memory usage of the Kudu Tserver process exceeds the threshold.

This alarm is cleared when the memory usage of the Tserver process becomes normal and the system considers that the Kudu instance service recovers.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29107	Critical	Yes

## Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Users cannot use the Kudu service.

## Possible Causes

The memory usage of the KuduTserver instance is too high.

## Handling Procedure

**Handle the Kudu instance exceptions.**

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the displayed page, locate alarm **ALM-29107 Tserver Process Memory Usage Exceeds the Threshold** and view the alarm source.

**Step 2** Choose **O&M > Alarm > Thresholds > Kudu**. Locate the alarm threshold, compare the memory usage of the KuduTserver instance in the cluster with the threshold, and find the node whose memory usage exceeds the threshold.

Add nodes or reschedule jobs to reduce the memory usage of the Tserver node or change the threshold.


**Step 3** Choose **O&M > Alarm** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [4](#).

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 5** Expand the **Service** drop-down list, and select **Kudu** for the target cluster.

**Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 7** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.244 ALM-38000 Kafka Service Unavailable

### Description

The system checks the Kafka service status every 30 seconds. This alarm is generated when the Kafka service is unavailable.

This alarm is cleared when the Kafka service recovers.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38000	Critical	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

## Impact on the System

The cluster cannot provide the Kafka service, and users cannot perform new Kafka tasks.

## Possible Causes

- The KrbServer service is abnormal.(Skip this step if the normal mode is used.)
- The ZooKeeper service is abnormal or does not respond.
- The Broker instance in the Kafka cluster are abnormal.

## Procedure

**Check the status of the KrbServer service. (Skip this step if the normal mode is used.)**

**Step 1** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **KrbServer**.

**Step 2** Check whether the running status of the KrbServer service is **Normal**.

- If yes, go to [Step 5](#).
- If no, go to [Step 3](#).

**Step 3** Rectify the fault by following the steps provided in **ALM-25500 KrbServer Service Unavailable**.

**Step 4** Perform [Step 2](#) again.

**Check the status of the ZooKeeper cluster.**

- Step 5** Check whether the running status of the ZooKeeper service is **Normal**.
- If yes, go to **Step 8**.
  - If no, go to **Step 6**.
- Step 6** If ZooKeeper service is stopped, start it, else rectify the fault by following the steps provided in **ALM-13000 ZooKeeper Service Unavailable**.
- Step 7** Perform **Step 5** again.

#### Check the Broker status.


- Step 8** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance** to go to the Kafka instances page.
- Step 9** Check whether all instances in **Roles** are running properly.
- If yes, go to **Step 11**.
  - If no, go to **Step 10**.
- Step 10** Select all Broker instances, choose **More** > **Restart Instance**, and check whether the instances restart successfully.

#### NOTE

During the restart of the Broker instance, if the current Topic is a single copy and is on the current Broker node, the Kafka service will be interrupted. Otherwise, the Kafka service will not be affected.

- If yes, go to **Step 11**.
  - If no, go to **Step 13**.
- Step 11** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** to check whether the running status is **Normal**.
- If yes, go to **Step 12**.
  - If no, go to **Step 13**.
- Step 12** Wait for 30 seconds and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 13**.

#### Collecting Fault Information

- Step 13** On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.
- Step 14** Select **Kafka** in the required cluster from the **Service** drop-down list.
- Step 15** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 16** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.245 ALM-38001 Insufficient Kafka Disk Capacity

### Description

The system checks the Kafka disk usage every 60 seconds and compares the actual disk usage with the threshold. The disk usage has a default threshold. This alarm is generated when the disk usage is greater than the threshold.

You can change the threshold in **O&M > Alarm > Thresholds**. Under the service list, choose **Kafka > Disk > Broker Disk Usage (Broker)** and change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the Kafka disk usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the Kafka disk usage is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38001	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition where the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

Kafka data write operations are affected.

## Possible Causes

- The configuration (such as number and size) of the disks for storing Kafka data cannot meet the requirement of the current service traffic, due to which the disk usage reaches the upper limit.
- Data retention time is too long, due to which the data disk usage reaches the upper limit.
- The service plan does not distribute data evenly, due to which the usage of some disks reaches the upper limit.

## Procedure

### Check the disk configuration of Kafka data.

**Step 1** On the FusionInsight Manager portal and click **O&M > Alarm > Alarms**.

**Step 2** In the alarm list, locate the alarm and obtain **HostName** from **Location**.

**Step 3** Click **Cluster > Name of the desired cluster > Hosts**.

**Step 4** In the host list, click the host name obtained in [Step 2](#).

**Step 5** Check whether the **Disk** area contains the partition name in the alarm.

- If yes, go to [Step 6](#).
- If no, manually clear the alarm and no further operation is required.

**Step 6** Check whether the disk partition usage contained in the alarm reaches 100% in the **Disk** area.

- If yes, handle the alarm by following the instructions in [Related Information](#).
- If no, go to [Step 7](#).

### Check the Kafka data storage duration.

**Step 7** Choose **Cluster > Name of the desired cluster > Services > Kafka > Configurations**.

**Step 8** Check whether the value of parameter **disk.adapter.enable** is set to **true**.

- If yes, go to [Step 10](#).
- If no, go to [Step 9](#).

**Step 9** Set the value of **disk.adapter.enable** to **true**. Check whether the value of **adapter.topic.min.retention.hours** is properly set.

- If yes, go to [Step 10](#).
- If no, adjust the data retention period based on service requirements.

**NOTICE**

If the disk auto-adaptation function is enabled, some historical data of specified topics is deleted. If the retention period of some topics cannot be adjusted, click **All Configurations** and add the topics to the value of the **disk.adapter.topic.blacklist** parameter.

**Step 10** Wait 10 minutes and check whether the usage of faulty disks reduces.

- If yes, wait until the alarm is cleared.
- If no, go to **Step 11**.

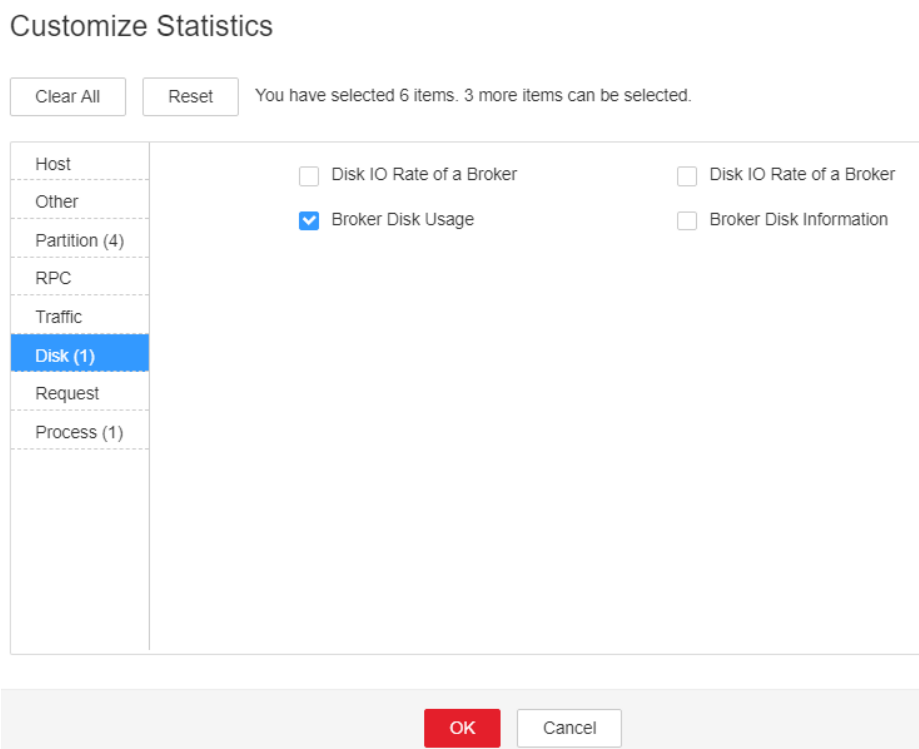
**Check the Kafka data plan.**

**Step 11** In the **Instance** area, click **Broker**. In the **Real Time** area of Broker, Click the drop-down menu in the Chart area and choose **Customize** to customize monitoring items.

**Step 12** In the dialog box, select **Disk > Broker Disk Usage** and click **OK**.

The Kafka disk usage information is displayed.

**Figure 7-131** Broker Disk Usage



**Step 13** View the information in **Step 12** to check whether there is only the disk parathion for which the alarm is generated in **Step 2**.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

**Step 14** Perform disk planning and mount a new disk again. Go to the **Instance Configurations** page of the node for which the alarm is generated, modify **log.dirs**, add other disk directories, and restart the Kafka instance.



 NOTE

During the restart of the Broker instance, if the current Topic is a single copy and is on the current Broker node, the Kafka service will be interrupted. Otherwise, the Kafka service will not be affected.

**Step 15** Determine whether to shorten the data retention time configured on Kafka based on service requirements and service traffic.

- If yes, go to [Step 16](#).
- If no, go to [Step 17](#).

**Step 16** Log in to FusionInsight Manager, select **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations**, and click **All Configurations**. In the search box on the right, enter **log.retention.hours**. The value of the parameter indicates the default data retention time of the topic. You can change the value to a smaller one.

 NOTE

- For a topic whose data retention time is configured alone, the modification of the data retention time on the Kafka Service Configuration page does not take effect.
- To modify the data retention time for a topic, use the Kafka client command-line interface (CLI) to configure the topic.

Example: **kafka-topics.sh --zookeeper "ZooKeeper IP address:2181/kafka" --alter --topic "Topic name" --config retention.ms= "retention time"**

**Step 17** Check whether the usage of some disks reaches the upper limit due to unreasonable configuration of the partitions of some topics. For example, the number of partitions configured for a topic with large data volume is smaller than the number of disks. In this case, the data is not evenly allocated to disks.

 NOTE

If you do not know which topics have a large amount of service data, perform the following steps:

1. Log in to an instance node based on the host node information obtained in [Step 2](#).
2. Go to the data directory (directory specified by **log.dirs** before the modification in [Step 14](#)).
3. Run the following command to check whether there is topic with partition that use large disk space.

```
du -h --max-depth=1 ./
```

- If yes, go to [Step 18](#).
- If no, go to [Step 19](#).

**Step 18** In the Kafka client CLI, run the following command to perform partition capacity expansion for the topic:

```
kafka-topics.sh --zookeeper "ZooKeeper IP address:2181/kafka" --alter --topic "Topic name" --partitions="New number of partitions"
```

 NOTE

- You are advised to set the new number of partitions to a multiple of the number of Kafka data disks.
- The step may not quickly clear the alarm, and you need to modify the data retention time in [Step 11](#) to gradually balance data allocation.

**Step 19** Determine whether to perform capacity expansion.

 **NOTE**

You are advised to perform capacity expansion for Kafka when the current disk usage exceeds 80%.

- If yes, go to [Step 20](#).
- If no, go to [Step 21](#).

**Step 20** Expand the disk capacity and check whether the alarm is cleared after capacity expansion.

- If yes, no further action is required.
- If no, go to [Step 22](#).


**Step 21** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 22](#).

**Collect fault information.**

**Step 22** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 23** Select **Kafka** in the required cluster from the **Service** drop-down list.

**Step 24** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 25** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

**Step 1** Log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Kafka > Instance**, stop the Broker instance whose status is **Restoring**, record the management IP address of the node where the Broker instance is located, and record **broker.id**. The value can be obtained by using the following method: Click the role name. On the **Configurations** page, select **All Configurations**, and search for the **broker.id** parameter.

**Step 2** Log in to the recorded management IP address as user **root**, and run the **df -lh** command to view the mounted directory whose disk usage is 100%, for example, **`\${BIGDATA\_DATA\_HOME}/kafka/data1**.

**Step 3** Go to the directory, run the **du -sh \*** command to view the size of each file in the directory, check whether files other than **kafka-logs** exist, and determine whether these files can be deleted or migrated.

- If yes, go to [Step 8](#).
- If no, go to [Step 4](#).

- Step 4** Go to the **kafka-logs** directory, run the **du -sh \*** command, select a partition folder to be moved. The naming rule is **Topic name-Partition ID**. Record the topic and partition.
- Step 5** Modify the **recovery-point-offset-checkpoint** and **replication-offset-checkpoint** files in the **kafka-logs** directory in the same way.
1. Decrease the number in the second line in the file. (To remove multiple directories, the number deducted is equal to the number of files to be removed.)
  2. Delete the line of the to-be-removed partition. (The line structure is "Topic name Partition ID Offset". Save the data before deletion. Subsequently, the content must be added to the file of the same name in the destination directory.)
- Step 6** Modify the **recovery-point-offset-checkpoint** and **replication-offset-checkpoint** files in the destination data directory. For example, **`\${BIGDATA\_DATA\_HOME}/kafka/data2/kafka-logs** in the same way.
- Increase the number in the second line in the file. (To move multiple directories, the number added is equal to the number of files to be moved.)
  - Add the to-be moved partition to the end of the file. (The line structure is "Topic name Partition ID Offset". You can copy the line data saved in [Step 5](#).)
- Step 7** Move the partition to the destination directory. After the partition is moved, run the **chown omm:wheel -R Partition directory** command to modify the directory owner group for the partition.
- Step 8** Log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > Kafka > Instance** to start the Broker instance.
- Step 9** Wait for 5 to 10 minutes and check whether the health status of the Broker instance is **Normal**.
- If yes, resolve the disk capacity insufficiency problem according to the handling method of "ALM-38001 Insufficient Kafka Disk Space" after the alarm is cleared.
  - If no, contact the O&M personnel.
- End

## 7.12.246 ALM-38002 Kafka Heap Memory Usage Exceeds the Threshold

### Description

The system checks the Kafka service status every 30 seconds. The alarm is generated when the heap memory usage of a Kafka instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times.

When the **Trigger Count** is 1, this alarm is cleared when the heap memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the heap memory usage is less than or equal to 90% of the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38002	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the available direct memory of the Kafka service is insufficient, a memory overflow occurs and the broker instance breaks down. As a result, the broker cannot provide read and write services properly.

## Possible Causes

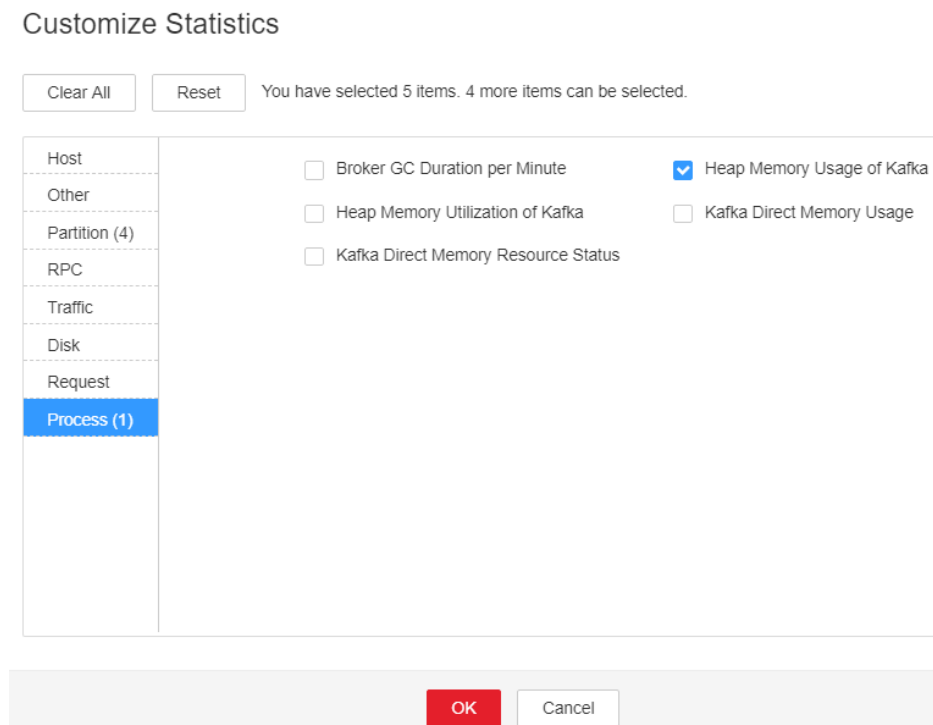
The heap memory of the Kafka instance is overused or the heap memory is inappropriately allocated.

## Procedure

### Check heap memory usage.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Kafka Heap Memory Usage Exceeds the Threshold > Location**. Check the host name of the instance involved in this alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Kafka > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down list in the upper right corner of the chart area, choose **Customize > Process > Heap Memory Usage of Kafka**, and click **OK**.

**Figure 7-132** Heap Memory Usage of Kafka



**Step 3** Check whether the used heap memory of Kafka reaches 95% of the maximum heap memory specified for Kafka.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

**Check the heap memory size configured for Kafka.**

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Kafka > Configurations > All Configurations > Broker(Role) > Environment**. Increase the value of **KAFKA\_HEAP\_OPTS** by referring to the Note.

**Figure 7-133** KAFKA\_HEAP\_OPTS

Parameter	Value
KAFKA_HEAP_OPTS	-Xmx6G -Xms6G


**NOTE**

- It is recommended that **-Xmx** and **-Xms** be set to the same value.
- You are advised to view **Heap Memory Usage of Kafka** by referring to **Step 2**, and set the value of **KAFKA\_HEAP\_OPTS** to twice the value of **Heap Memory Used by Kafka**.

**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

**Collect fault information.**

- Step 6** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 7** Select **Kafka** in the required cluster from the **Service** drop-down list.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.247 ALM-38004 Kafka Direct Memory Usage Exceeds the Threshold

## Description

The system checks the direct memory usage of the Kafka service every 30 seconds. This alarm is generated when the direct memory usage of a Kafka instance exceeds the threshold (80% of the maximum memory) for 10 consecutive times.

When the **Trigger Count** is 1, this alarm is cleared when the direct memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the direct memory usage is less than or equal to 90% of the threshold.

## Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38004	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the available direct memory of the Kafka service is insufficient, a memory overflow occurs and the broker instance breaks down. As a result, the broker cannot provide read and write services properly.

## Possible Causes

The direct memory of the Kafka instance is overused or the direct memory is inappropriately allocated.

## Procedure

**Check the direct memory usage.**

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Kafka Direct Memory Usage Exceeds the Threshold > Location** to check the host name of the instance for which the alarm is generated.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Kafka > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the Chart area and choose **Customize > Process > Kafka Direct Memory Usage**, and click **OK**.

**Figure 7-134** Kafka Direct Memory Usage

## Customize Statistics

Clear All    Reset    You have selected 5 items. 4 more items can be selected.

Host	<input type="checkbox"/> Broker GC Duration per Minute	<input type="checkbox"/> Heap Memory Usage of Kafka
Other	<input type="checkbox"/> Heap Memory Utilization of Kafka	<input checked="" type="checkbox"/> Kafka Direct Memory Usage
Partition (4)	<input type="checkbox"/> Kafka Direct Memory Resource Status	
RPC		
Traffic		
Disk		
Request		
Process (1)		

OK    Cancel

**Step 3** Check whether the used direct memory of Kafka reaches 80% of the maximum direct memory specified for Kafka.

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

**Check the direct memory size configured for the Kafka.**

**Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations** > **All Configurations** > **Broker(Role)** > **Environment** to increase the value of **-Xmx** configured in the **KAFKA\_HEAP\_OPTS** parameter by referring to the Note.

**NOTE**

- It is recommended that **-Xmx** and **-Xms** be set to the same value.
- You are advised to view **Kafka Direct Memory Usage** by referring to [Step 2](#), and set the value of **KAFKA\_HEAP\_OPTS** to twice the value of **Direct Memory Used by Kafka**.

**Step 5** Save the configuration and restart the Kafka service.

**NOTE**


If rolling restart is performed and the current Topic has multiple copies, services are not affected. Otherwise, the Kafka service will be unavailable during the restart, and upper-layer services that depend on the Kafka service will be affected.

**Step 6** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**



- Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 8** Select **Kafka** in the required cluster from the **Service** drop-down list.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.248 ALM-38005 GC Duration of the Broker Process Exceeds the Threshold

### Description

The system checks the garbage collection (GC) duration of the Broker process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default) for 3 consecutive times.

When the **Trigger Count** is 1, this alarm is cleared when the GC duration is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the GC duration is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38005	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

A long GC duration of the Broker process may interrupt the services.

## Possible Causes

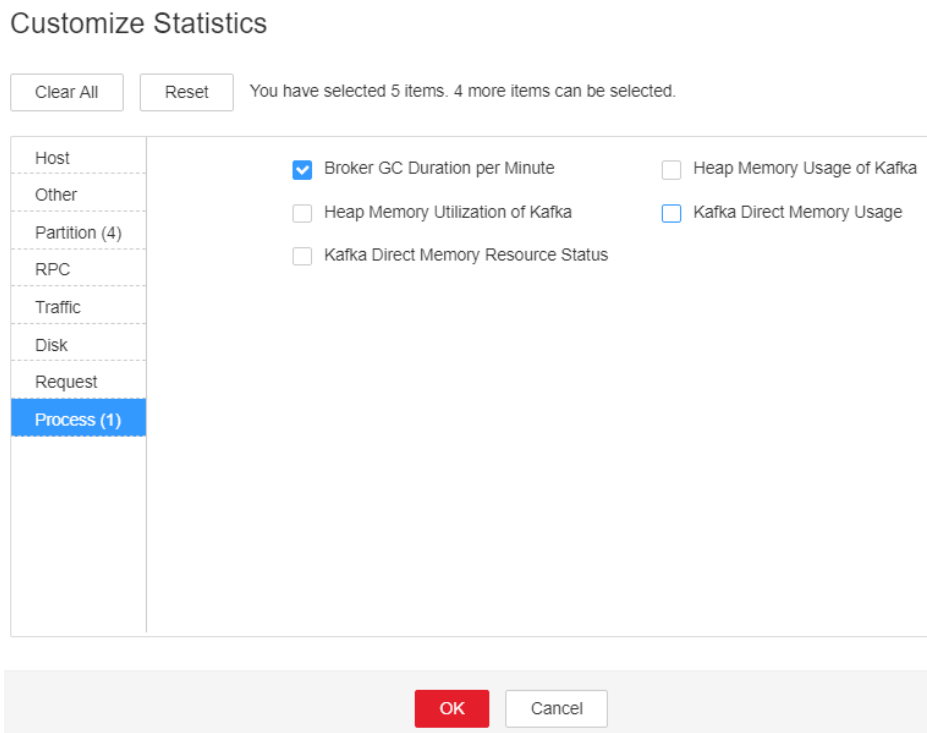
The Kafka GC duration of the node is too long or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

## Procedure

### Check the GC duration.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > GC Duration of the Broker Process Exceeds the Threshold > Location**. Check the host name of the instance involved in this alarm.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Name of the desired cluster > Services > Kafka > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down list in the upper right corner of the chart area, choose **Customize > Process > Broker GC Duration per Minute**, and click **OK**.

**Figure 7-135** Broker GC Duration per Minute



**Step 3** Check whether the GC duration of the Broker process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

**Check the direct memory size configured for the Kafka.**

**Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations** > **All Configurations** > **Broker(Role)** > **Environment** to increase the value of **-Xmx** configured in the **KAFKA\_HEAP\_OPTS** parameter by referring to the Note.

**NOTE**

- It is recommended that **-Xmx** and **-Xms** be set to the same value.
- You are advised to set the value of **KAFKA\_HEAP\_OPTS** to twice the value of **Direct Memory Used by Kafka**.

On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **Process** > **Kafka Direct Memory Resource Status** to check the value of **Direct Memory Used by Kafka**.

**Step 5** Save the configuration and restart the Kafka service.

**NOTE**


If rolling restart is performed and the current Topic has multiple copies, services are not affected. Otherwise, the Kafka service will be unavailable during the restart, and upper-layer services that depend on the Kafka service will be affected.

- Step 6** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **Kafka** in the required cluster from the **Service** drop-down list.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.249 ALM-38006 Percentage of Kafka Partitions That Are Not Completely Synchronized Exceeds the Threshold

### Description

The system checks the percentage of Kafka partitions that are not completely synchronized to the total number of partitions every 60 seconds. This alarm is generated when the percentage exceeds the threshold (50% by default) for 3 consecutive times.

When the **Trigger Count** is 1, this alarm is cleared when the percentage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the percentage is less than or equal to 90% of the threshold.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38006	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

Too many Kafka partitions that are not completely synchronized affect service reliability. In addition, data may be lost when leaders are switched.

## Possible Causes

Some nodes where the Broker instance resides are abnormal or stop running. As a result, replicas of some partitions in Kafka are out of the in-sync replicas (ISR) set.

## Procedure

### Check Broker instances.


- Step 1** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. The Kafka instances page is displayed.
- Step 2** Check whether faulty nodes exist among all Broker nodes.
  - If yes, record the host name of the node and go to **Step 3**.
  - If no, go to **Step 5**.
- Step 3** On the FusionInsight Manager portal, click **O&M** > **Alarm** > **Alarms** to check whether the fault described in **Step 2** exists in the alarm information and handle the alarm based on corresponding methods.
- Step 4** On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. The Kafka instances page is displayed.
- Step 5** Check whether stopped nodes exist among all Broker instance.
  - If yes, go to **Step 6**.
  - If no, go to **Step 7**.
- Step 6** Select all stopped Broker instances and click **Start Instance**.
- Step 7** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect fault information.**

**Step 8** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 9** Select **Kafka** in the required cluster from the **Service** drop-down list.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.250 ALM-38007 Status of Kafka Default User Is Abnormal

### Description

The system checks the default user of Kafka every 60 seconds. This alarm is generated when the system detects that the user status is abnormal.

**Trigger Count** is set to **1**. This alarm is cleared when the user status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38007	Critical	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host name for which the alarm is generated.
Trigger Condition	Specifies the condition that the Kafka default user status is abnormal.

## Impact on the System

If the Kafka default user status is abnormal, metadata synchronization between Brokers and interaction between Kafka and ZooKeeper will be affected, affecting service production, consumption, and topic creation and deletion.

## Possible Causes

- The Sssd service is abnormal.
- Some Broker instances stop running.

## Procedure

### Check whether the Sssd service is abnormal.

**Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms > Status of Kafka Default User Is Abnormal > Location** to check the host name of the instance for which the alarm is generated.

**Step 2** Find the host information in the alarm information and log in to the host.

**Step 3** Run the `id -Gn kafka` command and check whether "No such user" is displayed in the command output.

- If yes, record the host name of the node and go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On the FusionInsight Manager home page, choose **O&M > Alarm > Alarms**. Check whether there is **Sssd Service Exception** in the alarm information. If there is, handle the alarm based on alarm information.

### Check the running status of the Broker instance.

**Step 5** On the FusionInsight Manager home page, choose **Cluster > Name of the desired cluster > Services > Kafka > Instance**. The Kafka instance page is displayed.

**Step 6** Check whether there are stopped nodes on all Broker instances.

- If yes, go to [Step 7](#).
- If no, go to [Step 8](#).

**Step 7** Select all stopped Broker instances and click **Start Instance**.


**Step 8** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Collect fault information.**

**Step 9** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 10** In the **Service** area, select **Kafka** in the required cluster.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.251 ALM-38008 Abnormal Kafka Data Directory Status

### Description

The system checks the Kafka data directory status every 60 seconds. This alarm is generated when the system detects that the status of a data directory is abnormal.

**Trigger Count** is set to **1**. This alarm is cleared when the data directory status becomes normal.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38008	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.



Name	Meaning
HostName	Specifies the host name for which the alarm is generated.
DirName	Specifies the directory name for which the alarm is generated.
Trigger Condition	Specifies the condition that the Kafka data directory status is abnormal.

## Impact on the System

If the Kafka data directory status is abnormal, the current replicas of all partitions in the data directory are brought offline, and the data directory status of multiple nodes is abnormal at the same time. As a result, some partitions may become unavailable.

## Possible Causes

- The data directory permission is tampered with.
- The disk where the data directory is located is faulty.

## Procedure

### Check the permission on the faulty data directory.

**Step 1** Find the host information in the alarm information and log in to the host.

**Step 2** In the alarm information, check whether the data directory and its subdirectories belong to the omm:wheel group.

- If yes, record the host name of the node and go to [Step 4](#).
- If no, go to [Step 3](#).

**Step 3** Restore the owner group of the data directory and its subdirectories to omm:wheel.

- If yes, go to [Step 6](#).
- If no, go to [Step 5](#).

### Check whether the disk where the data directory is located is faulty.

**Step 4** In the upper-level directory of the data directory, create and delete files as user **omm**. Check whether data read/write on the disk is normal.

**Step 5** Replace or repair the disk where the data directory is located to ensure that data read/write on the disk is normal.

**Step 6** On the FusionInsight Manager home page, choose **Cluster > Services > Kafka > Instance**. On the Kafka instance page that is displayed, restart the Broker instance on the host recorded in [Step 2](#).

 NOTE

During the restart of the Broker instance, if the current Topic is a single copy and is on the current Broker node, the Kafka service will be interrupted. Otherwise, the Kafka service will not be affected.


**Step 7** After Broker is started, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect fault information.**

**Step 8** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 9** In the **Service** area, select **Kafka** in the required cluster.

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.252 ALM-38009 Busy Broker Disk I/Os (Applicable to Versions Later Than MRS 3.1.0)

 NOTE

- This section applies to versions later than MRS 3.1.0.
- If the alarm name is **ALM-38009 Kafka Topic Overload**, handle the alarm by following the instructions provided in [ALM-38009 Kafka Topic Overload \(Applicable to MRS 3.1.0 and Earlier Versions\)](#).

## Alarm Description

The system checks the I/O status of each Kafka disk every 60 seconds. This alarm is generated when the disk I/O of a Kafka data directory on a broker exceeds the threshold (80% by default).

Its **Trigger Count** is **3**. This alarm is cleared when the disk I/O is lower than the threshold (80% by default).

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
38009	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
DataDirectoryName	Specifies the name of the Kafka data directory with frequent disk I/Os.

## Impact on the System


The disk partition has frequent I/Os. Data may fail to be written to the Kafka topic for which the alarm is generated.

## Possible Causes

- There are many replicas configured for the topic.
- The parameter for batch writing producer's messages is inappropriately configured. The service traffic of this topic is too heavy, and the current partition configuration is inappropriate.

## Handling Procedure

**Check the number of topic replicas.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Locate the row that contains this alarm, click , and view the host name in **Location**.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services > Kafka > KafkaTopic Monitor**, search for the topic for which the alarm is generated, and check the number of replicas.
- Step 3** Reduce the replication factors of the topic (for example, reduce to **3**) if the number of replicas is greater than 3.

Run the following command on the FusionInsight client to replan the replicas of Kafka topics:

```
kafka-reassign-partitions.sh --zookeeper {zk_host}:{port}}kafka --reassignment-json-file {manual assignment json file path}--execute
```

For example:

```
/opt/client/Kafka/kafka/bin/kafka-reassign-partitions.sh --zookeeper 10.149.0.90:2181,10.149.0.91:2181,10.149.0.92:2181/kafka --reassignment-json-file expand-cluster-reassignment.json --execute
```

#### NOTE

In the **expand-cluster-reassignment.json** file, describe the brokers to which the partitions of the topic are migrated in the following format: {"partitions":[{"topic": "*topicName*", "partition": 1, "replicas": [1,2,3] }], "version":1}

**Step 4** Observe for a period of time and check whether the alarm is cleared. If the alarm persists, go to **Step 5**.

**Check the partition planning of the topic.**

**Step 5** On the **KafkaTopic Monitor** page, view **Topic Input Traffic** in the **Topic Traffic** area of each topic, obtain the topic with the largest value, and check the partitions of this topic as well as information about the host of these partitions.

**Step 6** Log in to the host queried in **Step 5** and run the **iostat -d -x** command to check the **%util** value of each disk.

```
:/opt/R3/FusionInsight_Manager/software/packs # iostat -d -x
Linux 3.0.76-0.11-default (189-39-172-162) 06/26/19 _x86_64_
Device: rrqm/s wrqm/s r/s w/s rsec/s wsec/s avgrq-sz avgqu-sz await svctm %util
xvda 0.04 44.44 1.26 21.94 43.62 531.02 24.78 0.03 1.44 0.56 1.30
xvde 0.16 431.84 13.78 82.51 284.32 4115.90 45.70 0.06 1.41 0.64 6.21
```

- If the **%util** value of each disk exceeds the threshold (**80%** by default), expand the Kafka disk capacity. After the capacity expansion, replan the topic partitions by referring to **Step 3**.
- If the **%util** values of the disks vary greatly, check the disk partition configuration of Kafka. For example, check the value of **log.dirs** in the **\$/BIGDATA\_HOME}/FusionInsight\_HD\_8.1.0.1/1\_14\_Broker/etc/server.properties** file.

Run the following command to view the **Filesystem** information:

```
df -h log.dirs value
```

The command output is as follows.

```
:/opt/R3/FusionInsight_Manager/software/packs # df -h /srv/BigData/kafka/data/kafka-logs/
filesystem Size Used Avail Use% Mounted on
/dev/xvda2 36G 21G 14G 62% /
:/opt/R3/FusionInsight_Manager/software/packs #
```

- If the partition where Filesystem is located matches the partition with a high **%util** value, plan Kafka partitions on idle disks, configure **log.dirs** as an idle disk directory, and replan topic partitions by referring to **Step 3**. Ensure that the partitions of the topic are evenly distributed to each disk.

**Step 7** Observe for a period of time and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, repeat **Step 5** to **Step 6** three times. Then, go to **Step 8**.

**Step 8** Observe for a period of time and check whether the alarm is cleared.


- If yes, no further action is required.

- If no, go to [Step 9](#).

#### Collect fault information.

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Expand the **Service** drop-down list, and select **Kafka** for the target cluster.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.253 ALM-38009 Kafka Topic Overload (Applicable to MRS 3.1.0 and Earlier Versions)

### NOTE

- This section applies to MRS 3.1.0 or earlier.
- If the alarm name is **ALM-38009 Busy Broker Disk I/Os**, handle the alarm by following the instructions provided in [ALM-38009 Busy Broker Disk I/Os \(Applicable to Versions Later Than MRS 3.1.0\)](#).

## Alarm Description

The system checks the overload status of each Kafka topic every 60 seconds. This alarm is generated when the percentage of partitions of a topic on the overloaded disk exceeds the threshold (40% by default).

Its **Trigger Count** is **1**. This alarm is cleared when the percentage of partitions of a topic on the overloaded disk is lower than the threshold (40% by default).

An overloaded disk refers to the disk whose I/O usage of a disk partition is greater than 80%.

For example:

The partitions of Topic A are distributed on three brokers. The I/O usages of the disk partitions on two brokers are greater than 80%.

The percentage of partitions on the overloaded disk is 2/3, greater than 40%, this alarm is generated.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
38009	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
TopicName	Specifies the Kafka topic for which the alarm was generated.

## Impact on the System


The disk partition has frequent I/Os. Data may fail to be written to the Kafka topic for which the alarm is generated.

## Possible Causes

- There are many replicas configured for the topic.
- The parameter for batch writing producer's messages is inappropriately configured. The service traffic of this topic is too heavy, and the current partition configuration is inappropriate.

## Handling Procedure

**Check the number of topic replicas.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. Locate the row that contains this alarm, click , and view the host name in **Location**.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services > Kafka > KafkaTopic Monitor**, search for the topic for which the alarm is generated, and check the number of replicas.
- Step 3** Reduce the replication factors of the topic (for example, reduce to **3**) if the number of replicas is greater than 3.

Run the following command on the FusionInsight client to replan the replicas of Kafka topics:

```
kafka-reassign-partitions.sh --zookeeper {zk_host}:{port}}kafka --
reassignment-json-file {manual assignment json file path}--execute
```

For example:

```
/opt/client/Kafka/kafka/bin/kafka-reassign-partitions.sh --zookeeper
10.149.0.90:2181,10.149.0.91:2181,10.149.0.92:2181/kafka --reassignment-json-
file expand-cluster-reassignment.json --execute
```

 NOTE

In the **expand-cluster-reassignment.json** file, describe the brokers to which the partitions of the topic are migrated in the following format: {"partitions":[{"topic": "*topicName*", "partition": 1, "replicas": [1,2,3] }], "version":1}

**Step 4** Observe for a period of time and check whether the alarm is cleared. If the alarm persists, go to **Step 5**.

**Check the partition planning of the topic.**

**Step 5** On the **KafkaTopic Monitor** page, view **Topic Input Traffic** in the **Topic Traffic** area of each topic, obtain the topic with the largest value, and check the partitions of this topic as well as information about the host of these partitions.

**Step 6** Log in to the host queried in **Step 5** and run the **iostat -d -x** command to check the **%util** value of each disk.

```
:/opt/R3/FusionInsight_Manager/software/packs # iostat -d -x
Linux 3.0.76-0.11-default (189-39-172-162) 06/26/19 _x86_64_
Device: rrqm/s wrqm/s r/s w/s rsec/s wsec/s avgrq-sz avgqu-sz await svctm %util
xvda 0.04 44.44 1.26 21.94 43.62 531.02 24.78 0.03 1.44 0.56 1.30
xvde 0.16 431.84 13.78 82.51 284.32 4115.90 45.70 0.06 1.41 0.64 6.21
```

- If the **%util** value of each disk exceeds the threshold (**80%** by default), expand the Kafka disk capacity. After the capacity expansion, replan the topic partitions by referring to **Step 3**.
- If the **%util** values of the disks vary greatly, check the disk partition configuration of Kafka. For example, check the value of **log.dirs** in the **\$/BIGDATA\_HOME}/FusionInsight\_HD\_8.1.0.1/1\_14\_Broker/etc/server.properties** file.

Run the following command to view the **Filesystem** information:

```
df -h log.dirs value
```

The command output is as follows.

```
:/opt/R3/FusionInsight_Manager/software/packs # df -h /srv/BigData/kafka/data/kafka-logs/
filesystem Size Used Avail Use% Mounted on
/dev/xvda2 36G 21G 14G 62% /
:/opt/R3/FusionInsight_Manager/software/packs #
```

- If the partition where Filesystem is located matches the partition with a high **%util** value, plan Kafka partitions on idle disks, configure **log.dirs** as an idle disk directory, and replan topic partitions by referring to **Step 3**. Ensure that the partitions of the topic are evenly distributed to each disk.

**Step 7** Observe for a period of time and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, repeat **Step 5** to **Step 6** three times. Then, go to **Step 8**.

**Step 8** Observe for a period of time and check whether the alarm is cleared.


- If yes, no further action is required.

- If no, go to [Step 9](#).

**Collect fault information.**

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Expand the **Service** drop-down list, and select **Kafka** for the target cluster.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.254 ALM-38010 Topics with Single Replica

### Description

The system checks the number of replicas of each topic every 60 seconds on the node where the Kafka Controller resides. This alarm is generated when there is one replica for a topic.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38010	Major	No

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.



Name	Meaning
TopicName	Specifies the list of topics for which the alarm is generated.

## Impact on the System

There is the single point of failure (SPOF) risk for topics with only one replica. When the node where the replica resides becomes abnormal, the partition does not have a leader, and services on the topic are affected.

## Possible Causes

- The number of replicas for the topic is incorrectly configured.

## Procedure

### Check the number of replicas for the topic.

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**, click  of this alarm, and view the **TopicName** list in **Location**.

**Step 2** Check whether replicas need to be added for the topic for which the alarm is generated.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

**Step 3** On the FusionInsight client, re-plan topic replicas and describe the partition distribution of the topic in the **add-replicas-reassignment.json** file in the following format: `{"partitions":[{"topic": "topic name", "partition": 1, "replicas": [1,2] }], "version":1}`. Then, run the following command to add replicas:

```
kafka-reassign-partitions.sh --zookeeper {zk_host}:{port}/kafka --reassignment-json-file {manual assignment json file path} --execute
```

For example:

```
/opt/client/Kafka/kafka/bin/kafka-reassign-partitions.sh --zookeeper 192.168.0.90:2181,192.168.0.91:2181,192.168.0.92:2181/kafka --reassignment-json-file add-replicas-reassignment.json --execute
```

**Step 4** Run the following command to check the task execution progress:

```
kafka-reassign-partitions.sh --zookeeper {zk_host}:{port}/kafka --reassignment-json-file {manual assignment json file path} --verify
```

For example:

```
/opt/client/Kafka/kafka/bin/kafka-reassign-partitions.sh --zookeeper 192.168.0.90:2181,192.168.0.91:2181,192.168.0.92:2181/kafka --reassignment-json-file add-replicas-reassignment.json --verify
```

**Step 5** After completing the handling operations or confirming that the alarm has no impact, manually clear the alarm on FusionInsight Manager.


**Step 6** After a period of time, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 7](#).

**Collect fault information.**

**Step 7** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 8** In the **Service** area, select **Kafka** in the required cluster.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

If the alarm has no impact, manually clear the alarm.

## Related Information

None

## 7.12.255 ALM-38011 User Connection Usage on Broker Exceeds the Threshold

### Description

The system checks the number of connections of each user on Broker every 30 seconds. This alarm is generated when the connection usage of a user on the Broker exceeds the threshold (80% by default) for 5 consecutive times.

The number of times that smoothing is performed is **5**. This alarm is cleared when the connection usage of a user on the Broker is less than the threshold.

The alarm can be automatically cleared. However, if the number of connections of a user suddenly becomes **0** and no connection is created, the alarm cannot be automatically cleared. You need to manually clear it.

### Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38011	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
UserName	Specifies the username for which the alarm is generated.

## Impact on the System

If the number of connections of a user is excessive, the user cannot create new connections to the Broker.

## Possible Causes

- The number of connections (created by a user) used by the client exceeds the preset threshold.
- The threshold for the connection usage does not meet service requirements.

## Procedure

**Check the number of connections established by the same user on the client.**

- Step 1** On the FusionInsight Manager home page, choose **O&M > Alarm > Alarms > User Connection Usage on Broker Exceeds the Threshold**. Check the host name and username of the Broker instance for which the alarm is generated in **Location**.
- Step 2** On FusionInsight Manager home page, choose **Cluster > Name of the desired cluster > Services > Kafka > Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down list in the upper right corner of the chart area, choose **Customize > Other**, and select **User Connection Usage on Broker, Maximum Number of User Connections on Broker, and Number of User Connections on Broker** to view the number of the current user connections on the Broker.
- Step 3** Observe the number of real-time connections of the current alarm user and check whether the real-time monitoring data of the current user exists.
  - If yes, go to **Step 4**.
  - If no, the current user has disconnected all connections. You need to clear the alarm manually, and no further action is required.

 NOTE

After the alarm user disconnects all connections, the monitoring data of the user disappears. In this case, the alarm will not be automatically cleared. You need to manually clear it.

**Step 4** Check whether the user is authorized by the service side.

If yes, go to [Step 7](#).

If no, go to [Step 5](#).

**Step 5** Run the following command on the client to limit the number of connections of the user. There are two configuration rules based on the following commands:

1. For the specific Broker and user, run the following command:

```
kafka-configs.sh --bootstrap-server <broker ip:port> --alter --add-config 'max.connections.per.user.overrides=[<username>:<connection.number>]' --entity-type brokers --entity-name <broker.id> --command-config Kafka/kafka/config/producer.properties
```

 NOTE

For unauthorized users, confirm with the service side to reduce the maximum number of connections of an unauthorized user or set the maximum number of connections to 0.

In the command, you need to specify the IP address and port number of Broker, set values of configuration items, and specify the **brokerId** and **username**. Here, the user refers to the authorized Kerberos user.

The configuration updated using the command line tool can take effect dynamically. The configuration becomes invalid after the service is restarted. To make the configuration take effect after the restart, choose **Cluster > Name of the desired cluster > Services > Kafka > Configurations > All Configurations > Broker > Server** on the FusionInsight Manager home page and update the configuration to **max.connections.per.user.overrides**.

2. For the specific use and default Broker (that is, all Broker instances in the cluster), run the following command:

```
kafka-configs.sh --bootstrap-server <broker ip:port> --alter --add-config 'max.connections.per.user.overrides=[<username>:<connection.number>]' --entity-type brokers ---entity-default --command-config Kafka/kafka/config/client.properties
```

Example:

```
kafka-configs.sh --bootstrap-server 10.153.3.26:21007 --alter --add-config 'max.connections.per.user.overrides=[showcase:4]' --entity-type brokers --entity-name 1 --command-config Kafka/kafka/config/client.properties
```

**Step 6** Check whether the maximum number of connections is 0 and whether the number of connections of the current user decreases or remains unchanged according to [Step 2](#).

- If yes, manually clear the alarm and no further action is required.
- If no, go to [Step 7](#).

**Step 7** Check whether the number of real-time connections and connection usage of the current user are sharply increased when they are compared with historical data, and whether have exceeded the specified maximum number of connections.

- If yes, go to [Step 8](#).
- If no, go to [Step 9](#).

 **NOTE**

If there is an obvious increase after the comparison and the maximum number of connections has reached the preset value, the connections of the user may be abnormal. You need to confirm with the service party.

**Check whether the number of user connections meets service requirements.**

**Step 8** Check whether the number of connections of the user meets service requirements.

- If yes, go to [Step 9](#).
- If no, contact the service party to rectify the fault.

 **NOTE**

If the number of user connections is abnormal, contact the service party to rectify the fault from the following aspects:

- Check whether new services are added so that the number of user connections increases sharply.
- Check whether handle leakage occurs on the code at the service side.

**Step 9** Consider whether to increase the maximum number of connections of the user.

- If yes, go to [Step 10](#).
- If no, go to [Step 12](#).

**Step 10** Increase the maximum number of connections based on the service requirements. Set the number of connections of the user on the Kafka client. For details, see [Step 5](#).

**Step 11** Wait for several minutes and then check whether the alarm is automatically cleared.

- If yes, go to [Step 12](#).
- If no, go to [Step 2](#).

**Step 12** Determine whether to add the user to the whitelist based on service requirements on the service side.

- If yes, go to [Step 13](#).
- If no, go to [Step 15](#).


 **NOTE**

To add a user to the whitelist, you need to restart the Kafka service. However, this operation will cause service interruption and affect service running. Therefore, you must confirm with the service side before performing this operation.

**Step 13** On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations** > **All Configurations** > **Broker(Role)** > **Server** to add the user to the **max.connections.per.user.whitelist** configuration item.

**Step 14** Restart the service for the modification to take effect. In addition, you need to manually clear the alarm, and no further action is required.

**Collect the fault information.**

- Step 15** On the FusionInsight Manager homepage, choose **O&M > Log > Download**.
- Step 16** Expand the **Service** drop-down list, and select **Kafka** for the target cluster.
- Step 17** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 18** Contact the O&M personnel and send the collected fault logs.
- End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.256 ALM-43001 Spark2x Service Unavailable

### Alarm Description

The system checks the Spark2x service status every 300 seconds. This alarm is generated when the Spark2x service is unavailable.

This alarm is cleared when the Spark2x service recovers.

#### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
43001	Critical	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

The Spark tasks submitted by users fail to be executed.

## Possible Causes

- The KrbServer service is abnormal.
- The LdapServer service is abnormal.
- ZooKeeper is abnormal.
- HDFS is abnormal.
- Yarn is abnormal.
- The corresponding Hive service is abnormal.
- The Spark2x assembly package is abnormal.
- The NameNode memory is insufficient.
- The memory of the Spark process is insufficient.

## Handling Procedure

If the alarm is abnormal Spark2x assembly packet, the Spark packet is abnormal. Wait for about 10 minutes. The alarm is automatically cleared.

**Check whether service unavailability alarms exist in services that Spark2x depends on.**

**Step 1** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**.

**Step 2** Check whether the following alarms exist in the alarm list:

- ALM-25500 KrbServer Service Unavailable
- ALM-25000 LdapServer Service Unavailable
- ALM-13000 ZooKeeper Service Unavailable
- ALM-14000 HDFS Service Unavailable
- ALM-18000 Yarn Service Unavailable
- ALM-16004 Hive Service Unavailable
- If yes, go to [Step 3](#).
- If no, go to [Step 4](#).

**Step 3** Handle the alarms based on the troubleshooting methods provided in the alarm help.

After the alarm is cleared, wait a few minutes and check whether the alarm GuardianService Unavailable is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the NameNode memory is insufficient.**

**Step 4** Check whether the NameNode memory is insufficient.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** Restart the NameNode to release the memory. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Check whether the memory of the Spark process is insufficient.**

**Step 6** Check whether the memory of the Spark process is insufficient due to memory-related modifications.

- If yes, go to [Step 7](#).
- If no, go to [Step 8](#).

**Step 7** Ensure that the memory of the Spark process is sufficient or expand the cluster capacity. Then, check whether this alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 8](#).

**Collect fault information.**

**Step 8** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 9** In the **Service** area, select the following nodes of the desired cluster. (Hive is the specific Hive service determined based on **ServiceName** in the alarm location information).

- KrbServer
- LdapServer
- ZooKeeper
- HDFS
- Yarn
- Hive

**Step 10** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 11** Contact O&M personnel and provide the collected logs.

----End



## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.257 ALM-43006 Heap Memory Usage of the JobHistory2x Process Exceeds the Threshold

## Description

The system checks the JobHistory2x Process status every 30 seconds. The alarm is generated when the heap memory usage of a JobHistory2x Process exceeds the threshold (95% of the maximum memory).

### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
43006	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the direct memory usage of the JobHistory2x process is too high, the performance deteriorates, and the process even becomes unavailable due to memory overflow. When it is unavailable, execution records of Spark tasks cannot be queried.

## Possible Causes

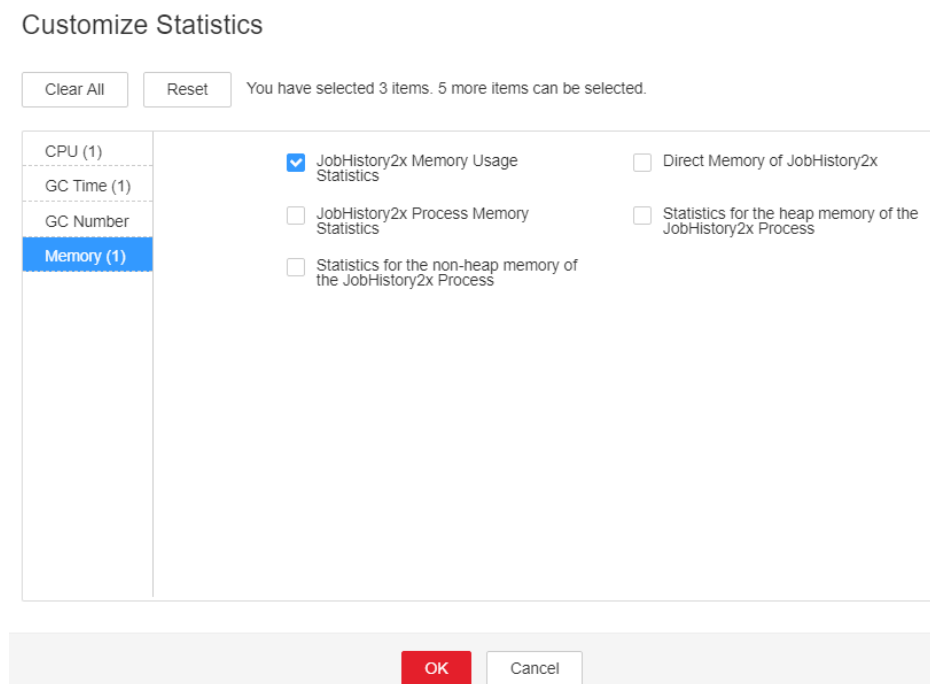
The heap memory of the JobHistory2x Process is overused or the heap memory is inappropriately allocated.

## Procedure

### Check heap memory usage.

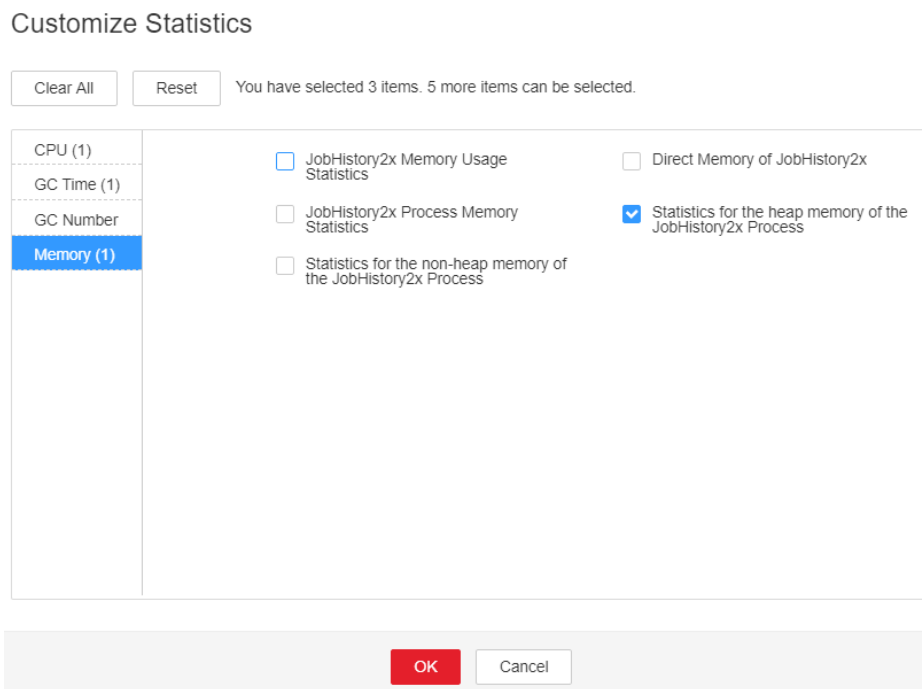
- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43006**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Instance** and click the JobHistory2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > Memory > JobHistory2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used heap memory of the JobHistory2x Process reaches the threshold(default value is 95%) of the maximum heap memory specified for JobHistory2x.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 7](#).

**Figure 7-136** JobHistory2x Memory Usage Statistics



**Step 3** On the FusionInsight Manager home page, choose **Cluster > Services > Spark2x > Instance**. Click **JobHistory2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize > Memory > Statistics for the heap memory of the JobHistory2x Process**, and click **OK**. Based on the alarm generation time, check the values of the used heap memory of the JobHistory2x process in the corresponding period and obtain the maximum value.

**Figure 7-137** Statistics for the heap memory of the JobHistory2x Process



**Step 4** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Configurations**, and click **All Configurations**. Choose **JobHistory2x > Default**. The default value of **SPARK\_DAEMON\_MEMORY** is 4GB. You can change the value according to the following rules: Ratio of the maximum heap memory usage of the JobHistory2x to the **Threshold** of the **JobHistory2x Heap Memory Usage Statistics (JobHistory2x)** in the alarm period. If this alarm is generated occasionally after the parameter value is adjusted, increase the value by 0.5 times. If the alarm is frequently reported after the parameter value is adjusted, increase the value by 1 time.

**NOTE**

On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Spark2x > Memory > JobHistory2x Heap Memory Usage Statistics (JobHistory2x)** to view **Threshold**.

**Step 5** Restart all JobHistory2x instances.

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.


**Step 6** After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **Spark2x** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.258 ALM-43007 Non-Heap Memory Usage of the JobHistory2x Process Exceeds the Threshold

## Description

The system checks the JobHistory2x Process status every 30 seconds. The alarm is generated when the non-heap memory usage of a JobHistory2x Process exceeds the threshold (95% of the maximum memory).

### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
43007	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the non-heap memory usage of the JobHistory2x process is too high, the performance deteriorates, and the process even becomes unavailable due to memory overflow. When it is unavailable, execution records of Spark tasks cannot be queried.

## Possible Causes

The non-heap memory of the JobHistory2x Process is overused or the non-heap memory is inappropriately allocated.

## Procedure

**Check non-heap memory usage.**

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43007**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Instance** and click the JobHistory2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > Memory > JobHistory2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used non-heap memory of the JobHistory2x Process reaches the threshold (default value is 95%) of the maximum non-heap memory specified for JobHistory2x.
  - If yes, go to **Step 3**.
  - If no, go to **Step 7**.

**Figure 7-138** JobHistory2x Memory Usage Statistics

Customize Statistics

Clear All    Reset    You have selected 3 items. 5 more items can be selected.

CPU (1)	<input checked="" type="checkbox"/> JobHistory2x Memory Usage Statistics	<input type="checkbox"/> Direct Memory of JobHistory2x
GC Time (1)	<input type="checkbox"/> JobHistory2x Process Memory Statistics	<input type="checkbox"/> Statistics for the heap memory of the JobHistory2x Process
GC Number	<input type="checkbox"/> Statistics for the non-heap memory of the JobHistory2x Process	
Memory (1)		

OK    Cancel

**Step 3** On the FusionInsight Manager home page, choose **Cluster > Services > Spark2x > Instance**. Click **JobHistory2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize > Memory > Statistics for the non-heap memory of the JobHistory2x Process**, and click **OK**. Based on the alarm generation time, check the values of the used non-heap memory of the JobHistory2x process in the corresponding period and obtain the maximum value.

**Figure 7-139** Statistics for the non-heap memory of the JobHistory2x Process

Customize Statistics

Clear All    Reset    You have selected 3 items. 5 more items can be selected.

CPU (1)	<input type="checkbox"/> JobHistory2x Memory Usage Statistics	<input type="checkbox"/> Direct Memory of JobHistory2x
GC Time (1)	<input type="checkbox"/> JobHistory2x Process Memory Statistics	<input type="checkbox"/> Statistics for the heap memory of the JobHistory2x Process
GC Number	<input checked="" type="checkbox"/> Statistics for the non-heap memory of the JobHistory2x Process	
Memory (1)		

OK    Cancel

**Step 4** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Configurations**, and click **All Configurations**. Choose **JobHistory2x > Default**. You can change the value of **-XX:MaxMetaspaceSize** in **SPARK\_DAEMON\_JAVA\_OPTS** according to the following rules: Ratio of the JobHistory2x non-heap memory usage to the **Threshold** of **JobHistory2x Non-Heap Memory Usage Statistics (JobHistory2x)** in the alarm period.

 **NOTE**

On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Spark2x > Memory > JobHistory2x Non-Heap Memory Usage Statistics (JobHistory2x)** to view **Threshold**.

**Step 5** Restart all JobHistory2x instances.

---

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

---


**Step 6** After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **Spark2x** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.259 ALM-43008 The Direct Memory Usage of the JobHistory2x Process Exceeds the Threshold

### Description

The system checks the JobHistory2x Process status every 30 seconds. The alarm is generated when the direct memory usage of a JobHistory2x Process exceeds the threshold (95% of the maximum memory).

 **NOTE**

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
43008	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the direct memory usage of the JobHistory process is too high, the performance deteriorates, and the process even becomes unavailable due to memory overflow. When it is unavailable, execution records of Spark tasks cannot be queried.

## Possible Causes

The direct memory of the JobHistory2x Process is overused or the direct memory is inappropriately allocated.

## Procedure

### Check direct memory usage.

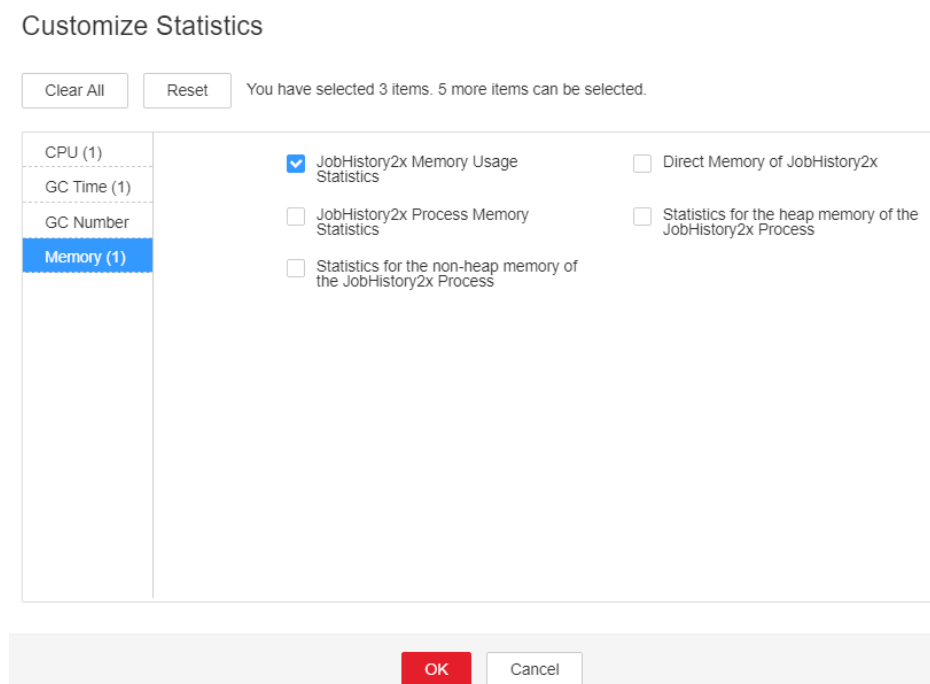
- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43008**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.



**Step 2** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Instance** and click the JobHistory2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > Memory > JobHistory2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used direct memory of the JobHistory2x Process reaches the threshold (default value is 95%) of the maximum direct memory specified for JobHistory2x.

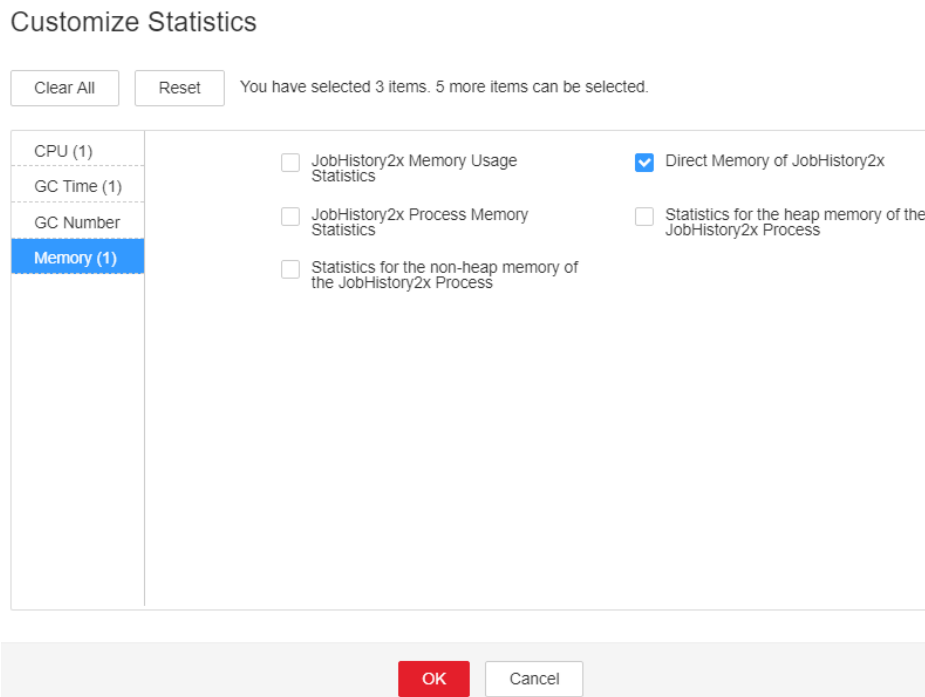
- If yes, go to **Step 3**.
- If no, go to **Step 7**.

**Figure 7-140** JobHistory2x Memory Usage Statistics



**Step 3** On the FusionInsight Manager home page, choose **Cluster > Services > Spark2x > Instance**. Click **JobHistory2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize > Memory > Direct Memory of JobHistory2x**, and click **OK**. Based on the alarm generation time, check the values of the used direct memory of the JobHistory2x process in the corresponding period and obtain the maximum value.

**Figure 7-141** Direct Memory of JobHistory2x



**Step 4** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Configurations**, and click **All Configurations**. Choose **JobHistory2x > Default**. The default value of **-XX:MaxDirectMemorySize** in **SPARK\_DAEMON\_JAVA\_OPTS** is 512 MB. You can change the value according to the following rules: Ratio of the maximum direct memory usage of the JobHistory2x to the **Threshold** of the **JobHistory2x Direct Memory Usage Statistics (JobHistory2x)** in the alarm period. If this alarm is generated occasionally after the parameter value is adjusted, increase the value by 0.5 times. If the alarm is frequently reported after the parameter value is adjusted, increase the value by 1 time. It is recommended that the value be less than or equal to the value of **SPARK\_DAEMON\_MEMORY**.

**NOTE**

On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Spark2x > Memory > JobHistory2x Direct Memory Usage Statistics (JobHistory2x)** to view **Threshold**.

**Step 5** Restart all JobHistory2x instances.


**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

**Step 6** After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

**Collect fault information.**

- Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.
- Step 8** Select **Spark2x** in the required cluster from the **Service**.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.260 ALM-43009 JobHistory2x Process GC Time Exceeds the Threshold

## Description

The system checks the garbage collection (GC) time of the JobHistory2x Process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (exceeds 5 seconds for three consecutive checks.) To change the threshold, choose **O&M > Alarm > Thresholds > Spark2x > GC Time > Total GC time in milliseconds (JobHistory2x)**. This alarm is cleared when the JobHistory2x GC time is shorter than or equal to the threshold.

### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
43009	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

The process performance deteriorates, and the process can even be unavailable. Historical execution records of Spark tasks cannot be queried.

## Possible Causes

The memory of JobHistory2x is overused, the heap memory is inappropriately allocated. As a result, GCs occur frequently.

## Procedure

### Check the GC time.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43009**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Instance** and click the JobHistory2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > GC Time > Garbage Collection (GC) Time of JobHistory2x** from the drop-down list box in the upper right corner and click **OK** to check whether the GC time is longer than the threshold(default value: 12 seconds).
  - If yes, go to **Step 3**.
  - If no, go to **Step 6**.

**Figure 7-142** Garbage Collection (GC) Time of JobHistory2x

## Customize Statistics

Clear All    Reset    You have selected 5 items. 3 more items can be selected.

CPU (1)	<input type="checkbox"/>	Garbage Collection (GC) Time of JobHistory2x
GC Time (1)	<input checked="" type="checkbox"/>	
GC Number	<input type="checkbox"/>	
Memory (3)	<input type="checkbox"/>	

OK    Cancel

**Step 3** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Configurations**, and click **All Configurations**. Choose **JobHistory2x > Default**. The default value of **SPARK\_DAEMON\_MEMORY** is 4GB. You can change the value according to the following rules: If this alarm is generated occasionally, increase the value by 0.5 times. If the alarm is frequently reported, increase the value by 1 time.

**Step 4** Restart all JobHistory2x instances.

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.


**Step 5** After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect fault information.**

**Step 6** On the FusionInsight Manager interface of active and standby clusters, choose **O&M > Log > Download**.

**Step 7** Select **Spark2x** in the required cluster from the **Service**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

# 7.12.261 ALM-43010 Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold

## Description

The system checks the JDBCServer2x Process status every 30 seconds. The alarm is generated when the heap memory usage of a JDBCServer2x Process exceeds the threshold (95% of the maximum memory).

### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
43010	Major	Yes

## Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.

Name	Meaning
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the heap memory usage of the JDBCServer2x process is too high, the performance deteriorates, and even memory overflow occurs. As a result, the JDBCServer2x process is unavailable, and Spark JDBC tasks are slow or fail to run.

## Possible Causes

The heap memory of the JDBCServer2x Process is overused or the heap memory is inappropriately allocated.

## Procedure

**Check heap memory usage.**

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43010**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Instance** and click the JDBCServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > Memory > JDBCServer2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used heap memory of the JDBCServer2x Process reaches the threshold (default value is 95%) of the maximum heap memory specified for JDBCServer2x.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 7](#).

**Figure 7-143** JDBCServer2x Memory Usage Statistics

Customize Statistics

Clear All    Reset    You have selected 6 items. 3 more items can be selected.

CPU (1)	<input checked="" type="checkbox"/> JDBCServer2x Memory Usage Statistics	<input type="checkbox"/> Direct Memory of JDBCServer2x
GC Time		
GC Number	<input type="checkbox"/> JDBCServer2x Process Memory Statistics	<input type="checkbox"/> Statistics for the heap memory of the JDBCServer2x Process
SQL (2)		
Session (2)		
Memory (1)	<input type="checkbox"/> Statistics for the non-heap memory of the JDBCServer2x Process	

OK    Cancel

**Step 3** On the FusionInsight Manager home page, choose **Cluster > Services > Spark2x > Instance**. Click **JDBCServer2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize > Memory > Statistics for the heap memory of the JDBCServer2x Process**, and click **OK**. Based on the alarm generation time, check the values of the used heap memory of the JDBCServer2x process in the corresponding period and obtain the maximum value.

**Figure 7-144** Statistics for the heap memory of the JDBCServer2x Process

Customize Statistics

Clear All    Reset    You have selected 6 items. 3 more items can be selected.

CPU (1)	<input type="checkbox"/> JDBCServer2x Memory Usage Statistics	<input type="checkbox"/> Direct Memory of JDBCServer2x
GC Time		
GC Number	<input type="checkbox"/> JDBCServer2x Process Memory Statistics	<input checked="" type="checkbox"/> Statistics for the heap memory of the JDBCServer2x Process
SQL (2)		
Session (2)		
Memory (1)	<input type="checkbox"/> Statistics for the non-heap memory of the JDBCServer2x Process	

OK    Cancel



**Step 4** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Configurations**, and click **All Configurations**. Choose **JDBCServer2x > Tuning**. The default value of **SPARK\_DRIVER\_MEMORY** is 4 GB. You can change the value according to the following rules: Ratio of the maximum heap memory usage of the JobHistory2x to the **Threshold** of the **JDBCServer2x Heap Memory Usage Statistics (JDBCServer2x)** in the alarm period. If this alarm is generated occasionally after the parameter value is adjusted, increase the value by 0.5 times. If the alarm is frequently reported after the parameter value is adjusted, increase the value by 1 time. In the case of large service volume and high concurrency, add instances.

 **NOTE**

On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Spark2x > Memory > JDBCServer2x Heap Memory Usage Statistics (JDBCServer2x)** to view **Threshold**.

**Step 5** Restart all JDBCServer2x instances.

---

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

---


**Step 6** After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **Spark2x** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.262 ALM-43011 Non-Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold

### Description

The system checks the JDBCServer2x Process status every 30 seconds. The alarm is generated when the non-heap memory usage of an JDBCServer2x Process exceeds the threshold (95% of the maximum memory).

#### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
43011	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

### Impact on the System

If the non-heap memory usage of the JDBCServer2x process is too high, the performance deteriorates, and even memory overflow occurs. As a result, the JDBCServer2x process is unavailable, and Spark JDBC tasks are slow or fail to run.

## Possible Causes

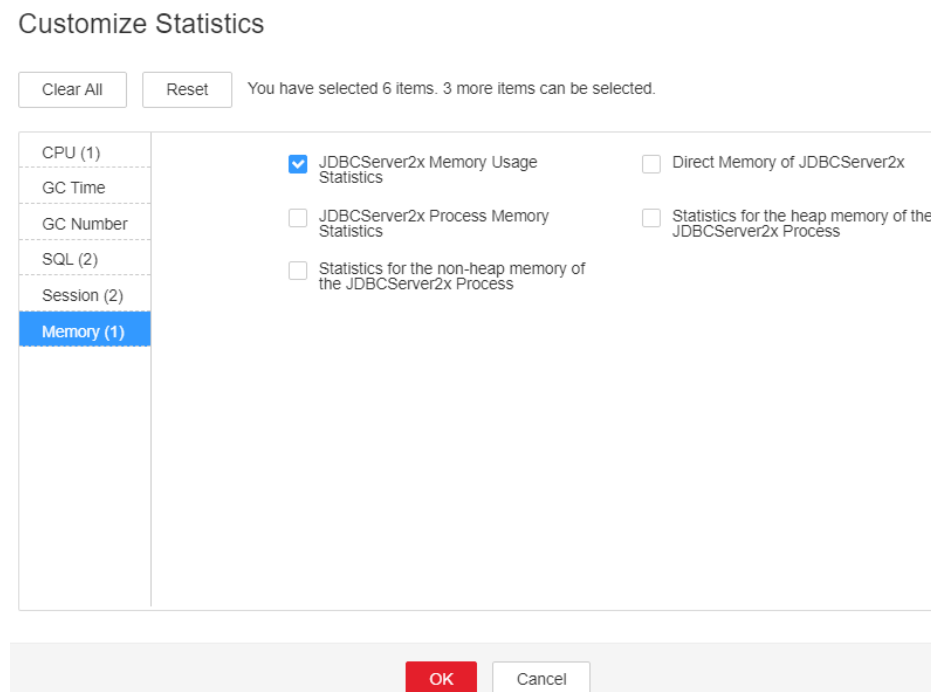
The non-heap memory of the JDBCServer2x Process is overused or the non-heap memory is inappropriately allocated.

## Procedure

**Check non-heap memory usage.**

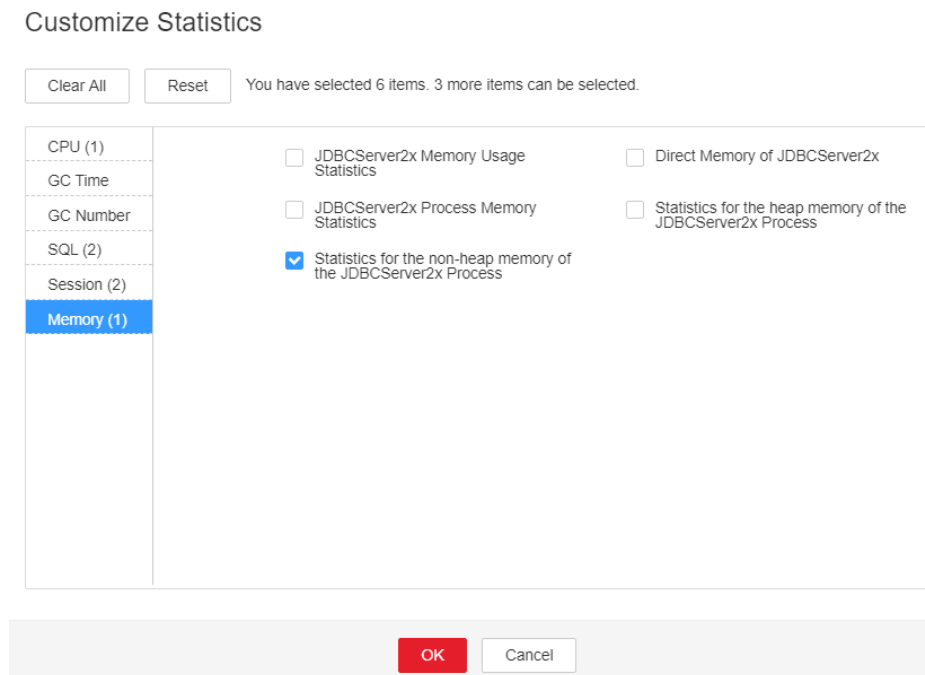
- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43011**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Instance** and click the JDBCServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > Memory > JDBCServer2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used non-heap memory of the JDBCServer2x Process reaches the threshold (default value is 95%) of the maximum non-heap memory specified for JDBCServer2x.
- If yes, go to **Step 3**.
  - If no, go to **Step 7**.

**Figure 7-145** JDBCServer2x Memory Usage Statistics



- Step 3** On the FusionInsight Manager home page, choose **Cluster > Services > Spark2x > Instance**. Click **JDBCServer2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize > Memory > Statistics for the non-heap memory of the JDBCServer2x Process**, and click **OK**. Based on the alarm generation time, check the values of the used non-heap memory of the JDBCServer2x process in the corresponding period and obtain the maximum value.

**Figure 7-146** Statistics for the non-heap memory of the JDBCServer2x Process



**Step 4** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Configurations**, and click **All Configurations**. Choose **JDBCServer2x > Tuning**. You can change the value of **-XX: MaxMetaspaceSize** in **spark.driver.extraJavaOptions** according to the following rules: Ratio of the JDBCServer2x non-heap memory usage to the **Threshold** of **JDBCServer2x Non-Heap Memory Usage Statistics ( JDBCServer2x)** in the alarm period.

**NOTE**

On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Spark2x > Memory > JDBCServer2x Non-Heap Memory Usage Statistics (JDBCServer2x)** to view **Threshold**.

**Step 5** Restart all JDBCServer2x instances.

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.


**Step 6** After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

**Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

**Step 8** Select **Spark2x** in the required cluster from the **Service**.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.263 ALM-43012 Direct Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold

### Description

The system checks the JDBCServer2x Process status every 30 seconds. The alarm is generated when the direct heap memory usage of a JDBCServer2x Process exceeds the threshold (95% of the maximum memory).

#### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
43012	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.

Name	Meaning
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

## Impact on the System

If the direct memory usage of the JDBCServer2x process is too high, the performance deteriorates, and even memory overflow occurs. As a result, the JDBCServer2x process is unavailable, and Spark JDBC tasks are slow or fail to run.

## Possible Causes

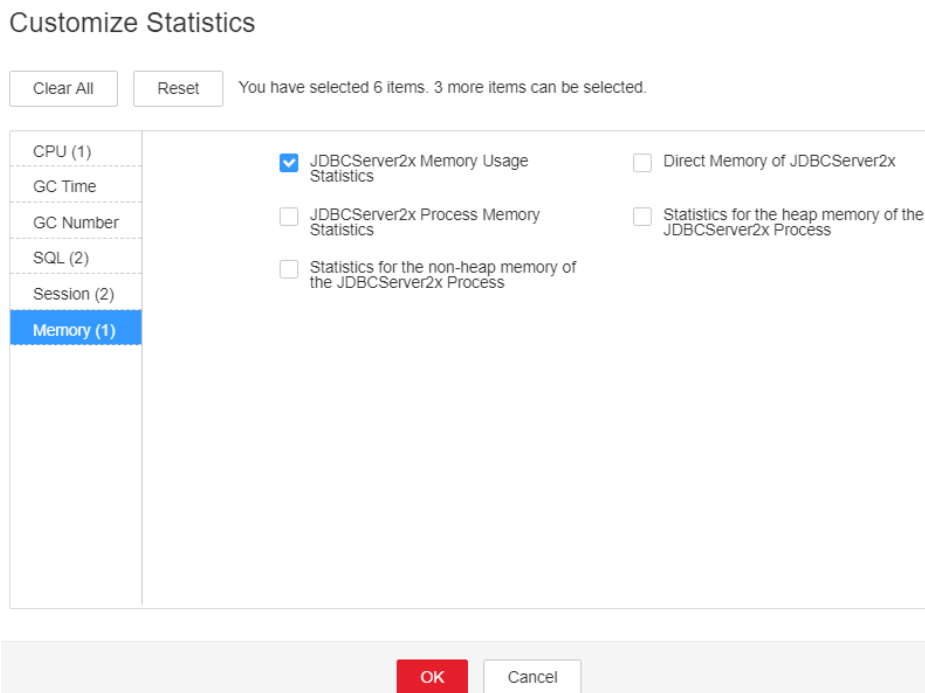
The direct heap memory of the JDBCServer2x Process is overused or the direct heap memory is inappropriately allocated.

## Procedure

### Check direct heap memory usage.

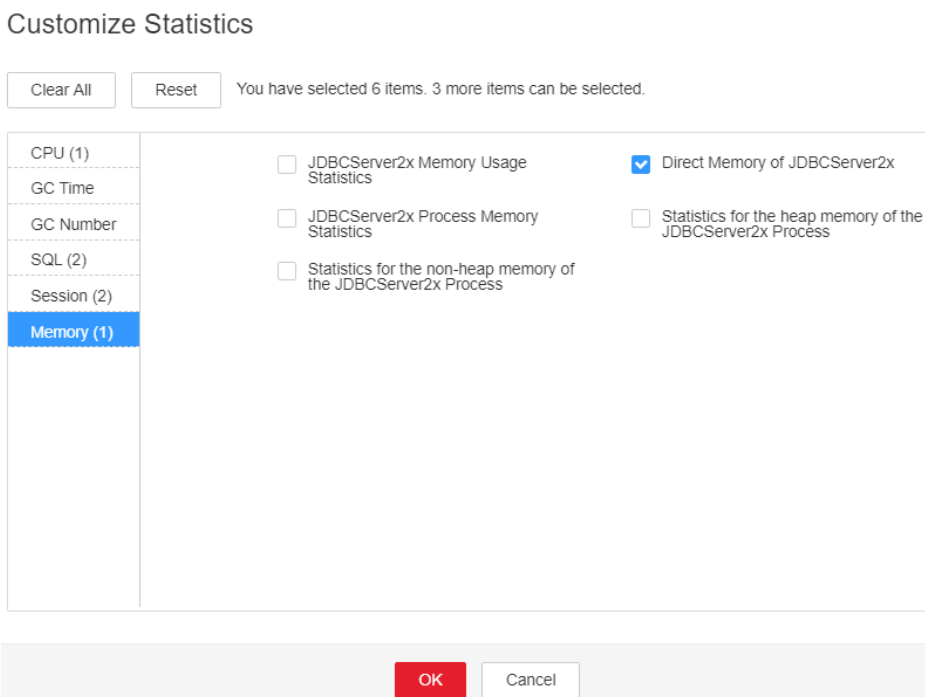
- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43012**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Instance** and click the JDBCServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > Memory > JDBCServer2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used direct heap memory of the JDBCServer2x Process reaches the threshold (default value is 95%) of the maximum direct heap memory specified for JDBCServer2x.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 7](#).

**Figure 7-147** JDBCServer2x Memory Usage Statistics



**Step 3** On the FusionInsight Manager home page, choose **Cluster > Services > Spark2x > Instance**. Click **JDBCServer2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize > Memory > Direct Memory of JDBCServer2x**, and click **OK**. Based on the alarm generation time, check the values of the used direct memory of the JDBCServer2x process in the corresponding period and obtain the maximum value.

**Figure 7-148** Direct Memory of JDBCServer2x



- Step 4** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Configurations**, and click **All Configurations**. Choose **JDBCServer2x > Tuning**. The default value of **-XX:MaxDirectMemorySize** in **spark.driver.extraJavaOptions** is 512 MB. You can change the value according to the following rules: Ratio of the maximum direct memory usage of the JDBCServer2x to the **Threshold** of the **JDBCServer2x Direct Memory Usage Statistics (JDBCServer2x)** in the alarm period. If this alarm is generated occasionally after the parameter value is adjusted, increase the value by 0.5 times. If the alarm is frequently reported after the parameter value is adjusted, increase the value by 1 time. In the case of large service volume and high service concurrency, you are advised to add instances.

 **NOTE**

On the FusionInsight Manager home page, choose **O&M > Alarm > Thresholds > Spark2x > Memory > JDBCServer2x Direct Memory Usage Statistics (JDBCServer2x)** to view **Threshold**.

- Step 5** Restart all JDBCServer2x instances.

---

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

---


- Step 6** After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Collect fault information.**

- Step 7** On the FusionInsight Manager portal, choose **O&M > Log > Download**.

- Step 8** Select **Spark2x** in the required cluster from the **Service**.

- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

- Step 10** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None



## 7.12.264 ALM-43013 JDBCServer2x Process GC Time Exceeds the Threshold

### Description

The system checks the garbage collection (GC) time of the JDBCServer2x Process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (exceeds 5 seconds for three consecutive checks.) To change the threshold, choose **O&M > Alarm > Thresholds > Spark2x > GC Time > Total GC time in milliseconds (JDBCServer2x)**. This alarm is cleared when the JDBCServer2x GC time is shorter than or equal to the threshold.

#### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
43013	Major	Yes

### Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

### Impact on the System

If the GC duration exceeds the threshold, the performance of the JDBCServer2x process deteriorates, and the process can even be unavailable. As a result, Spark JDBC tasks are slow or fail to run.

## Possible Causes

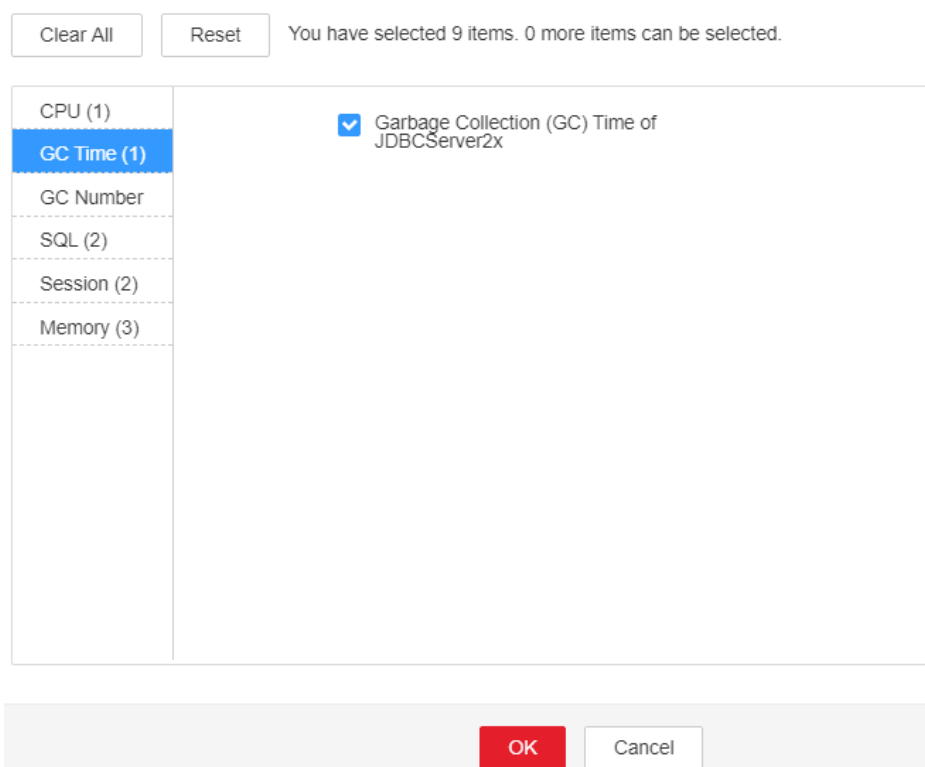
The memory of JDBCServer2x is overused, the heap memory is inappropriately allocated. As a result, GCs occur frequently.

## Procedure

### Check the GC time.

- Step 1** On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43013**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Instance** and click the JDBCServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > GC Time > Garbage Collection (GC) Time of JDBCServer2x** from the drop-down list box in the upper right corner and click **OK** to check whether the GC time is longer than the threshold(default value: 12 seconds).
- If yes, go to [Step 3](#).
  - If no, go to [Step 6](#).

**Figure 7-149** Garbage Collection (GC) Time of JDBCServer2x  
Customize Statistics



- Step 3** On the FusionInsight Manager portal, choose **Cluster > Services > Spark2x > Configurations**, and click **All Configurations**. Choose **JDBCServer2x > Default**. The default value of **SPARK\_DRIVER\_MEMORY** is 4 GB. If this alarm is generated occasionally, increase the value by 0.5 times. If the alarm is frequently reported,

increase the value by 1 time. In the case of large service volume and high service concurrency, you are advised to add instances.

**Step 4** Restart all JDBCServer2x instances.

---

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

---


**Step 5** After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect fault information.**

**Step 6** On the FusionInsight Manager interface of active and standby clusters, choose **O&M > Log > Download**.

**Step 7** Select **Spark2x** in the required cluster from the **Service**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Related Information

None

## 7.12.265 ALM-43017 JDBCServer2x Process Full GC Number Exceeds the Threshold

### Description

The system checks the number of Full garbage collection (GC) times of the JDBCServer2x process every 60 seconds. This alarm is generated when the detected Full GC number exceeds the threshold (exceeds 12 for three consecutive checks.) You can change the threshold by choosing **O&M > Alarm > Thresholds > Spark2x > GC number > Full GC Number of JDBCServer2x**. This alarm is cleared when the Full GC number of the JDBCServer2x process is less than or equal to the threshold.

 NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
43017	Major	Yes

## Parameters

Name	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the full GC times exceeds the threshold, the performance of the JDBCServer2x process deteriorates, and the process can even be unavailable. As a result, Spark JDBC tasks are slow or fail to run.

## Possible Causes

The heap memory usage of the JDBCServer2x process is excessively large, or the heap memory is inappropriately allocated. As a result, Full GC occurs frequently.

## Procedure

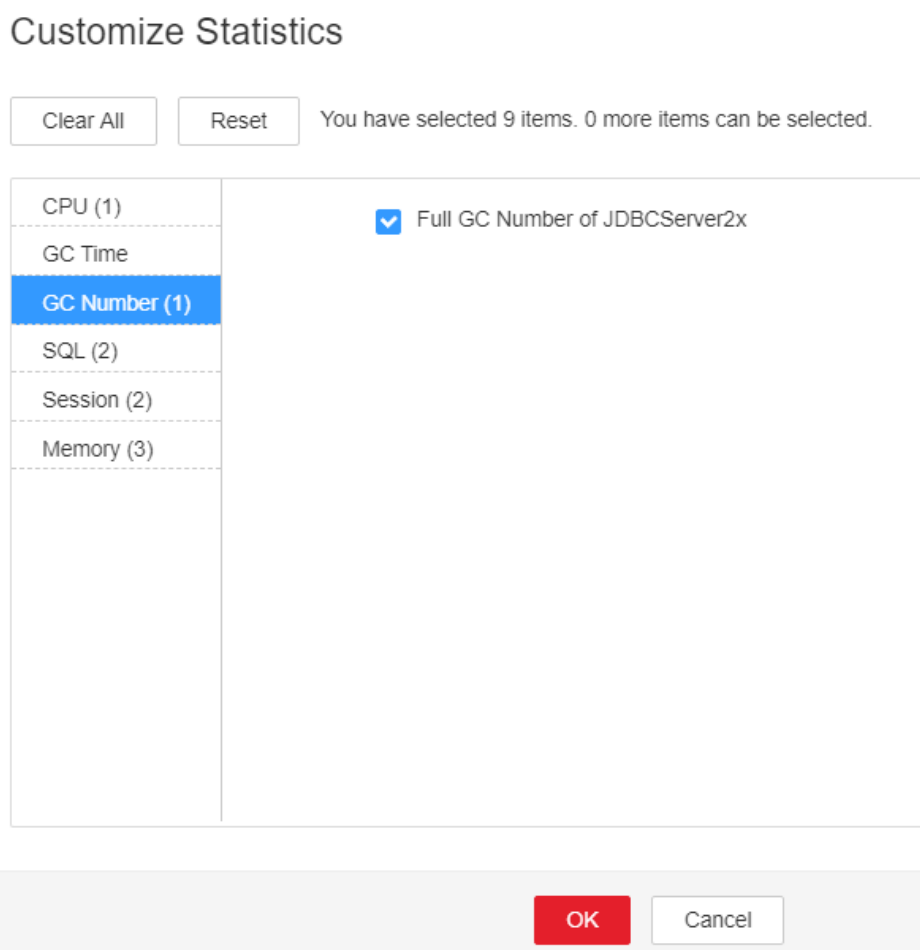
### Check the number of Full GCs.

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, select this alarm, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.

**Step 2** Choose **Cluster > Services > Spark2x > Instance**. On the displayed page, click the JDBCServer2x for which the alarm is reported. On the **Dashboard** page that is displayed, click the drop-down menu in the Chart area and choose **Customize > GC Number > Full GC Number of JDBCServer2x** in the upper right corner and click **OK**. Check whether the number of Full GCs of the JDBCServer2x process is greater than the threshold(default value: 12).

- If it is, go to **Step 3**.
- If it is not, go to **Step 6**.

**Figure 7-150** Full GC Number of JDBCServer2x



**Step 3** Choose **Cluster > Services > Spark2x > Configurations > All Configurations**. On the displayed page, choose **JDBCServer2x > Tuning**. The default value of **SPARK\_DRIVER\_MEMORY** is 4GB. You can change the value according to the following rules: If this alarm is generated occasionally, increase the value by 0.5 times. If the alarm is frequently reported, increase the value by 1 time. In the case of large service volume and high concurrency, add instances.

**Step 4** Restart all JDBCServer2x instances.

### NOTICE

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.


**Step 5** After 10 minutes, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to [Step 6](#).

**Collect fault information.**

**Step 6** Log in to FusionInsight Manager, and choose **O&M > Log > Download**.

**Step 7** Select **Spark2x** in the required cluster from the **Service** drop-down list.

**Step 8** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.266 ALM-43018 JobHistory2x Process Full GC Number Exceeds the Threshold

### Description

The system checks the number of Full garbage collection (GC) times of the JobHistory2x process every 60 seconds. This alarm is generated when the detected Full GC number exceeds the threshold (exceeds 12 for three consecutive checks.) You can change the threshold by choosing **O&M > Alarm > Thresholds > Spark2x > GC number > Full GC Number of JobHistory2x**. This alarm is cleared when the Full GC number of the JobHistory2x process is less than or equal to the threshold.

#### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
43018	Major	Yes

## Parameters

Name	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

The process performance deteriorates, and the process can even be unavailable. Historical execution records of Spark tasks cannot be queried.

## Possible Causes

The heap memory usage of the JobHistory2x process is excessively large, or the heap memory is inappropriately allocated. As a result, Full GC occurs frequently.

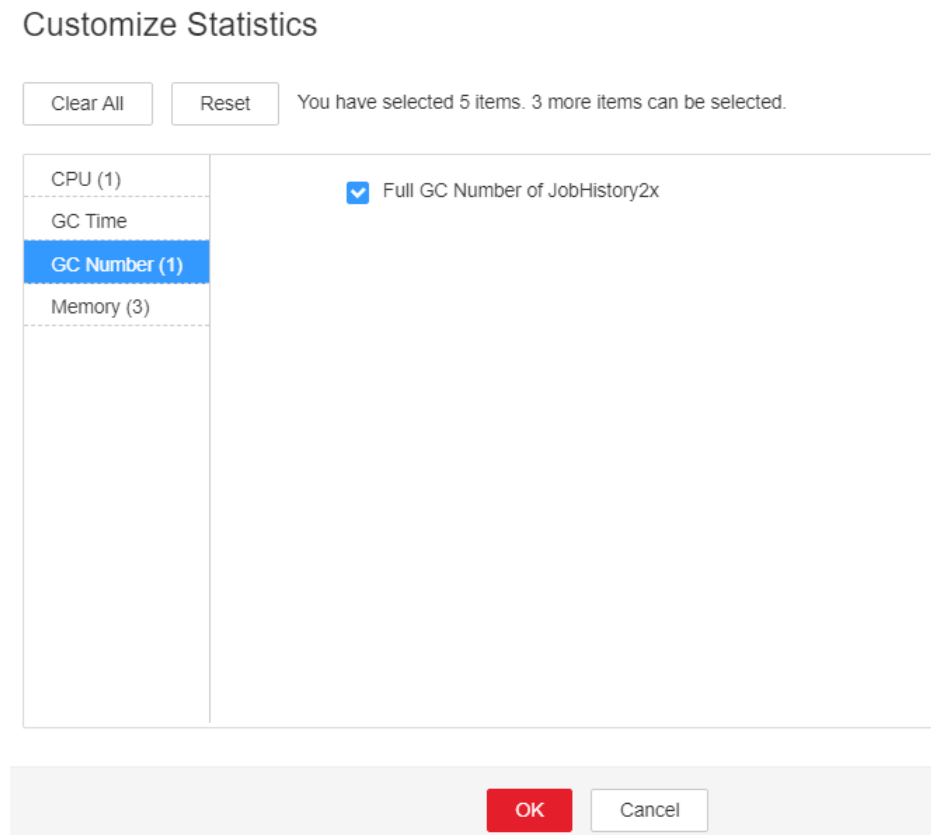
## Procedure

### Check the number of Full GCs.

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, select this alarm, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** Choose **Cluster > Services > Spark2x > Instance**. On the displayed page, click the JobHistory2x for which the alarm is reported. On the **Dashboard** page that is displayed, click the drop-down menu in the Chart area and choose **Customize > GC Number > Full GC Number of JobHistory2x** in the upper right corner and click **OK**. Check whether the number of Full GCs of the JobHistory2x process is greater than the threshold(default value: 12).
  - If it is, go to [Step 3](#).

- If it is not, go to [Step 6](#).

**Figure 7-151** Full GC Number of JobHistory2x



**Step 3** Choose **Cluster > Services > Spark2x > Configurations > All Configurations**. On the displayed page, choose **JobHistory2x > Default**. The default value of **SPARK\_DAEMON\_MEMORY** is 4GB. You can change the value according to the following rules: If this alarm is generated occasionally, increase the value by 0.5 times. If the alarm is frequently reported, increase the value by 1 time.

**Step 4** Restart all JobHistory2x instances.

#### NOTICE

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

**Step 5** After 10 minutes, check whether the alarm is cleared.


- If it is, no further action is required.
- If it is not, go to [Step 6](#).

#### Collect fault information.

**Step 6** Log in to FusionInsight Manager, and choose **O&M > Log > Download**.

**Step 7** Select **Spark2x** in the required cluster from the **Service**.



**Step 8** Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

**Step 9** Contact the O&M personnel and send the collected logs.

----End

## Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.267 ALM-43019 Heap Memory Usage of the IndexServer2x Process Exceeds the Threshold

## Description

The system checks the IndexServer2x process status every 30 seconds. The alarm is generated when the heap memory usage of a IndexServer2x process exceeds the threshold (95% of the maximum memory).

### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Severity	Auto Clear
43019	Major	Yes

## Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the heap memory usage of the IndexServer2x process is too high, the performance deteriorates, and even memory overflow occurs. As a result, the IndexServer2x process is unavailable, and Carbon tasks with indexing enabled are slow or fail to run.

## Possible Causes

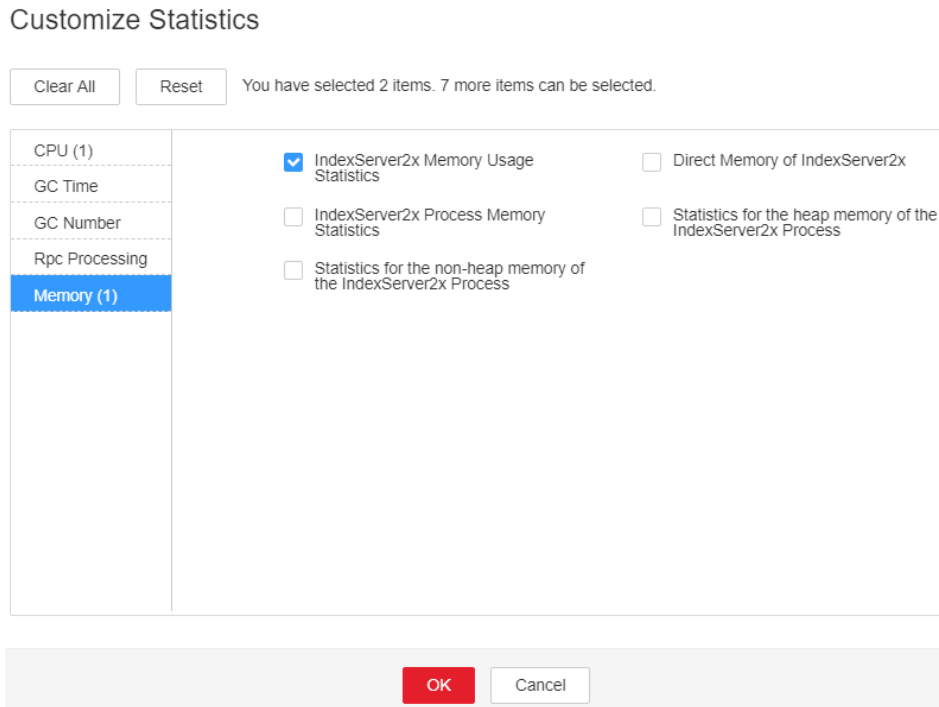
The heap memory of the IndexServer2x process is overused or the heap memory is inappropriately allocated.

## Procedure

**Check the heap memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the displayed alarm list, choose the alarm for which the ID is **43019**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize > Memory > IndexServer2x Memory Usage Statistics > OK**. Check whether the heap memory used by the IndexServer2x process reaches the maximum heap memory threshold (95% by default).
  - If the threshold is reached, go to **Step 3**.
  - If the threshold is not reached, go to **Step 7**.

**Figure 7-152** IndexServer2x Memory Usage Statistics



**Step 3** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize > Memory > Statistics for the heap memory of the IndexServer2x Process > OK**. Based on the alarm generation time, check the values of the used heap memory of the IndexServer2x process in the corresponding period and obtain the maximum value.

**Figure 7-153** Statistics for the heap memory of the IndexServer2x Process

Customize Statistics

You have selected 2 items. 7 more items can be selected.

CPU (1)	<input type="checkbox"/> IndexServer2x Memory Usage Statistics	<input type="checkbox"/> Direct Memory of IndexServer2x
GC Time	<input type="checkbox"/> IndexServer2x Process Memory Statistics	<input checked="" type="checkbox"/> Statistics for the heap memory of the IndexServer2x Process
GC Number	<input type="checkbox"/> Statistics for the non-heap memory of the IndexServer2x Process	
Rpc Processing		
Memory (1)		

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Configurations > All Configuration > IndexServer2x > Tuning**. The default value of the **SPARK\_DRIVER\_MEMORY** parameter is 4 GB. You can change the value based on the ratio of the maximum heap memory used by the IndexServer2x process to the threshold specified by **IndexServer2x Heap Memory Usage Statistics (IndexServer2x)** in the alarm period. If the alarm persists after the parameter value is changed, increase the value by 0.5 times. If the alarm is generated frequently, double the rate.

 **NOTE**

On FusionInsight Manager, you can choose **O&M > Alarm > Thresholds > Spark2x > Memory > IndexServer2x Heap Memory Usage Statistics (IndexServer2x)** to view the threshold.

**Step 5** Restart all IndexServer2x instances.

---

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

---


**Step 6** After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to [Step 7](#).

**Collect fault information.**

**Step 7** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 8** Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and provide the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Reference

None

## 7.12.268 ALM-43020 Non-Heap Memory Usage of the IndexServer2x Process Exceeds the Threshold

### Description

The system checks the IndexServer2x process status every 30 seconds. The alarm is generated when the non-heap memory usage of the IndexServer2x process exceeds the threshold (95% of the maximum memory).

#### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

### Attribute

Alarm ID	Severity	Auto Clear
43020	Major	Yes

### Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the non-heap memory usage of the IndexServer2x process is too high, the performance deteriorates, and even memory overflow occurs. As a result, the IndexServer2x process is unavailable, and Carbon tasks with indexing enabled are slow or fail to run.

## Possible Causes


The non-heap memory of the IndexServer2x process is overused or the non-heap memory is inappropriately allocated.

## Procedure

**Check non-heap memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the displayed alarm list, choose the alarm for which the ID is **43020**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize > Memory > IndexServer2x Memory Usage Statistics > OK**. Check whether the non-heap memory used by the IndexServer2x process reaches the maximum non-heap memory threshold (95% by default).
  - If the threshold is reached, go to [Step 3](#).
  - If the threshold is not reached, go to [Step 7](#).

**Figure 7-154** IndexServer2x Memory Usage Statistics



**Step 3** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize > Memory > Statistics for the non-heap memory of the IndexServer2x Process > OK**. Based on the alarm generation time, check the values of the used non-heap memory of the IndexServer2x process in the corresponding period and obtain the maximum value.

**Figure 7-155** Statistics for the non-heap memory of the IndexServer2x Process

Customize Statistics

Clear All    Reset    You have selected 2 items. 7 more items can be selected.

CPU (1)	<input type="checkbox"/> IndexServer2x Memory Usage Statistics	<input type="checkbox"/> Direct Memory of IndexServer2x
GC Time	<input type="checkbox"/> IndexServer2x Process Memory Statistics	<input type="checkbox"/> Statistics for the heap memory of the IndexServer2x Process
GC Number		
Rpc Processing		
Memory (1)	<input checked="" type="checkbox"/> Statistics for the non-heap memory of the IndexServer2x Process	

OK    Cancel

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Configurations > All Configurations > IndexServer2x > Tuning**. You can change the value of **XX:MaxMetaspaceSize** in the **spark.driver.extraJavaOptions** parameter based on the ratio of the maximum non-heap memory used by the IndexServer2x process to the threshold specified by **IndexServer2x Non-Heap Memory Usage Statistics (IndexServer2x)** in the alarm period.

**NOTE**

On FusionInsight Manager, you can choose **O&M > Alarm > Thresholds > Spark2x > Memory > IndexServer2x Non-Heap Memory Usage Statistics (IndexServer2x)** to view the threshold.

**Step 5** Restart all IndexServer2x instances.

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

**Step 6** After 10 minutes, check whether the alarm is cleared.


- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to [Step 7](#).

**Collect fault information.**

**Step 7** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 8** Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.



**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and provide the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Reference

None

## 7.12.269 ALM-43021 Direct Memory Usage of the IndexServer2x Process Exceeds the Threshold

### Description

The system checks the IndexServer2x process status every 30 seconds. The alarm is generated when the direct heap memory usage of a IndexServer2x process exceeds the threshold (95% of the maximum memory).

#### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

### Attribute

Alarm ID	Severity	Auto Clear
43021	Major	Yes

### Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the direct memory usage of the IndexServer2x process is too high, the performance deteriorates, and even memory overflow occurs. As a result, the IndexServer2x process is unavailable, and Carbon tasks with indexing enabled are slow or fail to run.

## Possible Causes

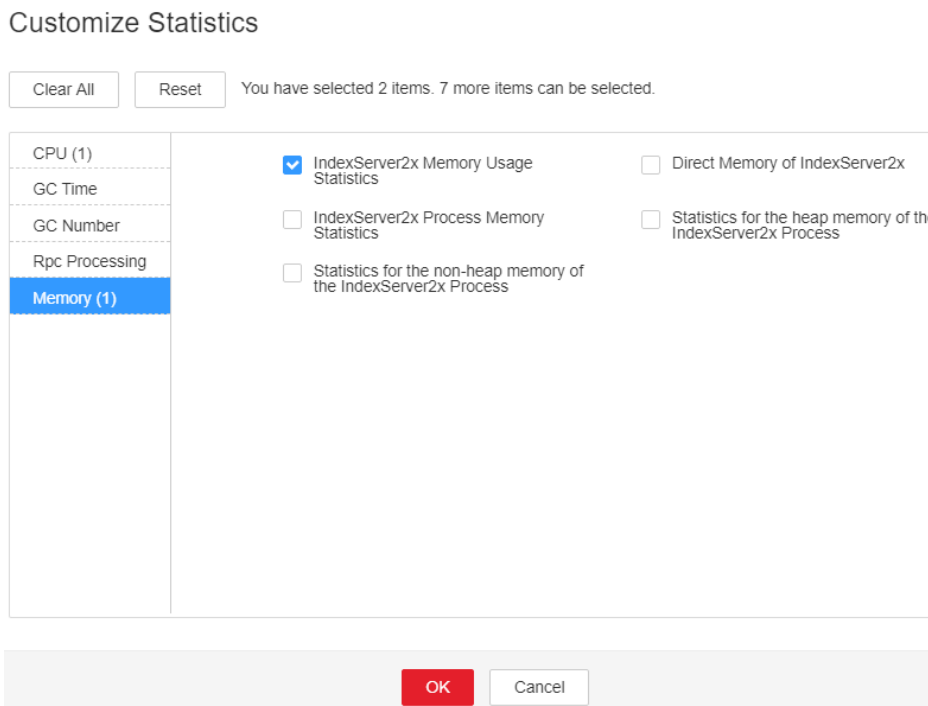
The direct heap memory of the IndexServer2x process is overused or the direct heap memory is inappropriately allocated.

## Procedure

**Check direct heap memory usage.**

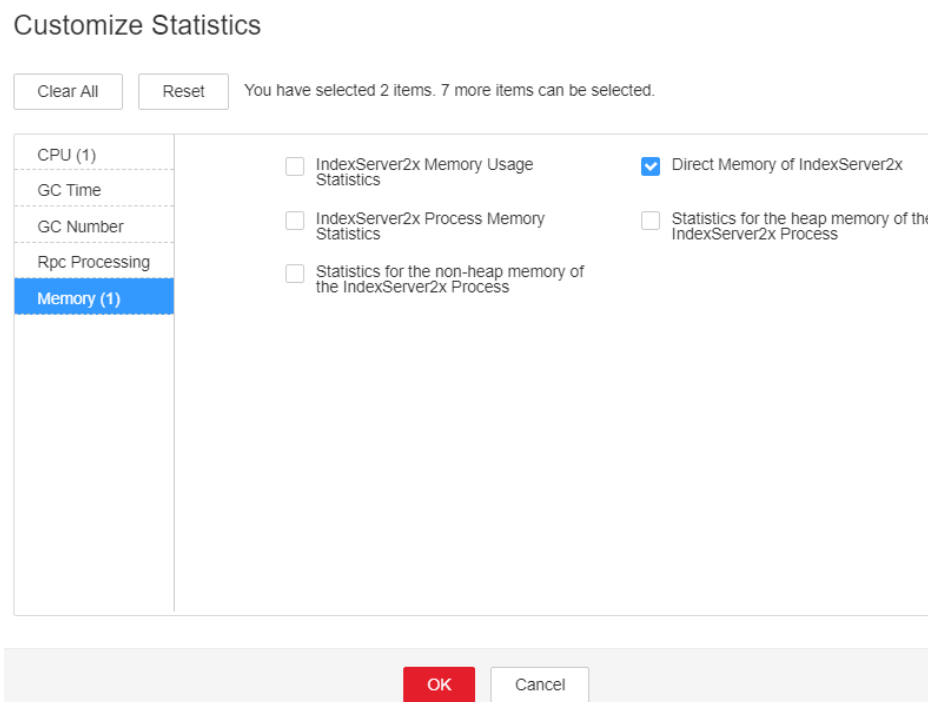
- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the displayed alarm list, choose the alarm for which the ID is **43021**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize > Memory > IndexServer2x Memory Usage Statistics > OK**. Check whether the direct memory used by the IndexServer2x process reaches the maximum direct memory threshold.
  - If the threshold is reached, go to [Step 3](#).
  - If the threshold is not reached, go to [Step 7](#).

**Figure 7-156** IndexServer2x Memory Usage Statistics



**Step 3** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize > Memory > Direct Memory of IndexServer2x > OK**. Based on the alarm generation time, check the values of the used direct memory of the IndexServer2x process in the corresponding period and obtain the maximum value.

**Figure 7-157** Direct Memory of IndexServer2x



**Step 4** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Configurations > All Configurations > IndexServer2x > Tuning**. You can change the value of **XX:MaxDirectMemorySize** (the default value is 512 MB) in the **spark.driver.extraJavaOptions** parameter based on the ratio of the maximum direct memory used by the IndexServer2x process to the threshold specified by **IndexServer2x Direct Memory Usage Statistics (IndexServer2x)** in the alarm period. If the alarm persists after the parameter value is changed, increase the value by 0.5 times. If the alarm is generated frequently, double the rate.

 **NOTE**

On FusionInsight Manager, you can choose **O&M > Alarm > Thresholds > Spark2x > Memory > IndexServer2x Direct Memory Usage Statistics (IndexServer2x)** to view the threshold.

**Step 5** Restart all IndexServer2x instances.

---

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

---


**Step 6** After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to [Step 7](#).

**Collect fault information.**

**Step 7** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 8** Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 10** Contact the O&M personnel and provide the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Reference

None

## 7.12.270 ALM-43022 IndexServer2x Process GC Time Exceeds the Threshold

### Description

The system checks the GC time of the IndexServer2x process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (12

seconds) for three consecutive times. To change the threshold, choose **O&M > Alarm > Thresholds > Spark2x > GC Time > Total GC time in milliseconds (IndexServer2x)**. This alarm is cleared when the IndexServer2x GC time is shorter than or equal to the threshold.

#### NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Severity	Auto Clear
43022	Major	Yes

## Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the GC duration exceeds the threshold, the performance of the IndexServer2x process deteriorates, and the process can even be unavailable. As a result, Carbon tasks with indexing enabled are slow or fail to run.

## Possible Causes

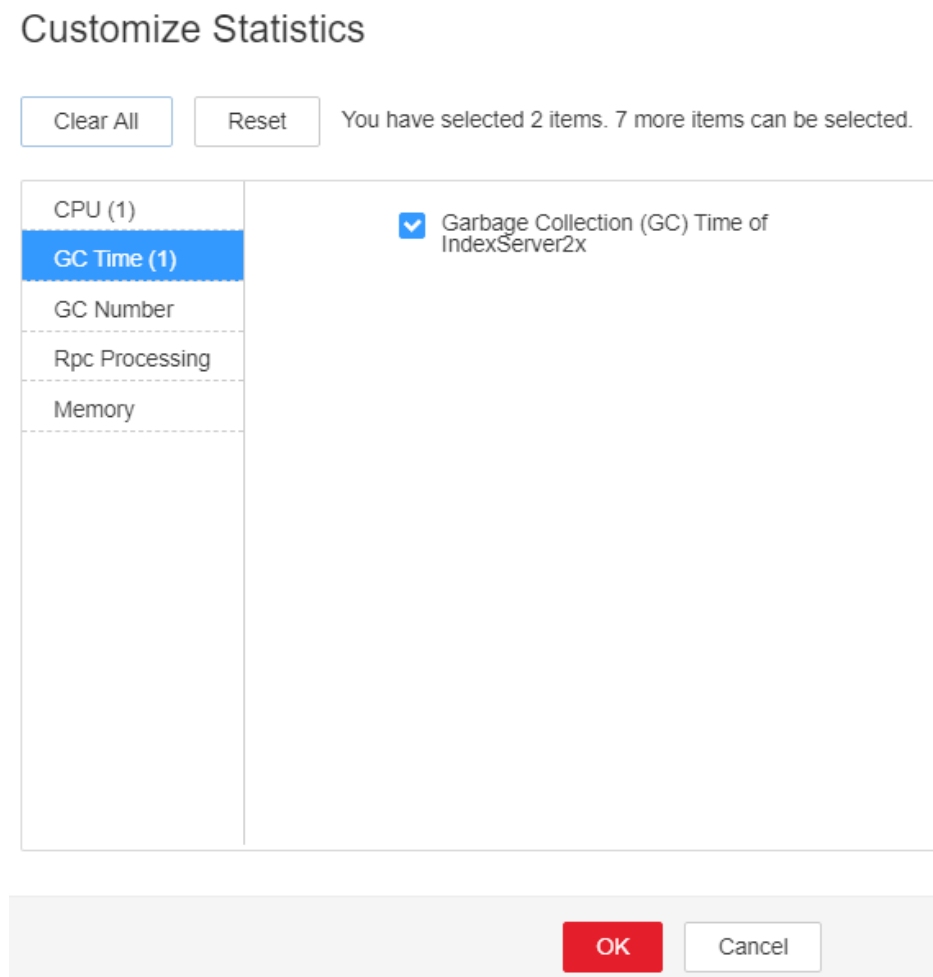
The heap memory of the IndexServer2x process is overused or the heap memory is inappropriately allocated. As a result, GC occurs frequently.

## Procedure

**Check the GC time.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the displayed alarm list, choose the alarm with ID **43022**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Instance** and click the **IndexServer2x** for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize > GC Time > Garbage Collection (GC) Time of IndexServer2x** from the drop-down list box in the upper right corner and click **OK** to check whether the GC time is longer than the threshold (default value: 12 seconds).
- If the threshold is reached, go to **Step 3**.
  - If the threshold is not reached, go to **Step 6**.

**Figure 7-158** Garbage Collection (GC) Time of IndexServer2x



- Step 3** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Configurations > All Configurations > IndexServer2x > Default**. The default value of the **SPARK\_DRIVER\_MEMORY** is 4 GB. You can change the value according to the following rules: Increase the value of the **SPARK\_DRIVER\_MEMORY** parameter 1.5 times to its default value. If this alarm is still generated occasionally after the adjustment, increase the value by 0.5 times. Double the value if the alarm is reported frequently.

**Step 4** Restart all IndexServer2x instances.

---

**NOTICE**

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

---


**Step 5** After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to [Step 6](#).

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 9** Contact the O&M personnel and provide the collected logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Reference

None

## 7.12.271 ALM-43023 IndexServer2x Process Full GC Number Exceeds the Threshold

### Description

The system checks the Full GC number of the IndexServer2x process every 60 seconds. This alarm is generated when the detected Full GC number exceeds the threshold (12) for three consecutive times. You can change the threshold by choosing **O&M > Alarm > Thresholds > Spark2x > GC Number > Full GC Number of IndexServer2x**. This alarm is cleared when the Full GC number of the IndexServer2x process is less than or equal to the threshold. This alarm is cleared when the Full GC number of the IndexServer2x process is less than or equal to the threshold.

 **NOTE**

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

## Attribute

Alarm ID	Severity	Auto Clear
43023	Major	Yes

## Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the GC times exceeds the threshold, the performance of the IndexServer2x process deteriorates, and the process can even be unavailable. As a result, Carbon tasks with indexing enabled are slow or fail to run.

## Possible Causes

The heap memory of the IndexServer2x process is overused or the heap memory is inappropriately allocated. As a result, Full GC occurs frequently.

## Procedure


### Check the number of Full GCs.

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the displayed alarm list, choose the alarm with the ID **43023**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Instance** and click the IndexServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the chart area and choose **Customize > GC Number > Full GC Number of IndexServer2x** from the drop-down list box in the upper right corner and click **OK** to check whether the GC number is larger than the threshold (default value: 12).
  - If the threshold is reached, go to [Step 3](#).



- If the threshold is not reached, go to [Step 6](#).

**Figure 7-159** Full GC Number of IndexServer2x



**Step 3** On FusionInsight Manager, choose **Cluster > Services > Spark2x > Configurations > All Configurations > IndexServer2x > Tuning**. The default value of the **SPARK\_DRIVER\_MEMORY** is 4 GB. You can change the value according to the following rules: If this alarm is generated occasionally, increase the value by 0.5 times. Double the value if the alarm is reported frequently. In the case of large service volume and high service concurrency, you are advised to add instances.

**Step 4** Restart all IndexServer2x instances.


#### NOTICE

When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

**Step 5** After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to [Step 6](#).

**Collect fault information.**

- Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 9** Contact the O&M personnel and send the collected fault logs.

----End

## Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

## Reference

None

## 7.12.272 ALM-44000 Presto Service Unavailable

### Alarm Description

The system checks the Presto service status every 60 seconds. This alarm is generated when the system detects that Presto is unavailable.

This alarm is cleared when the Presto service recovers.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
44000	Critical	Yes

### Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

Presto cannot run SQL queries.

## Possible Causes


- The Presto coordinator or worker process is faulty.
- The network communication between Presto coordinator and worker instances is interrupted.

## Handling Procedure

**Step 1** Check the status of the coordinator and worker processes.

1. Log in to FusionInsight Manager and choose **Cluster > Services > Presto**. On the page that is displayed, click the **Instance** tab. In the Presto instance list, check whether the status of all coordinator or worker instances is **Unknown**.
  - If yes, go to **2**.
  - If no, go to **1**.
2. In the upper part of the Presto instance list, choose **More > Restart Service** to restart the coordinator and worker processes.
3. In the alarm list, check whether ALM-44000 Presto Service Unavailable is cleared.
  - If yes, no further action is required.
  - If no, go to **1** in **Step 2**.

**Step 2** Collect fault information.

1. On FusionInsight Manager, choose **System > Export Log**.
2. Select **Presto** for **Service**.
3. Click  in the upper right corner.  
Set **Start Time** and **End Time** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **OK**.
4. Contact the O&M engineers and send the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## 7.12.273 ALM-44004 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold

### Alarm Description

This alarm is generated when the system detects that the number of queuing tasks in a resource group exceeds the threshold. The system queries the number of queuing tasks in a resource group through the JMX interface. You can choose **Components > Presto > Service Configuration** (switch **Basic** to **All**) > **Presto >**

**resource-groups** to configure a resource group. You can choose **Components > Presto > Service Configuration** (switch **Basic** to **All**) > **Coordinator > Customize > resourceGroupAlarm** to configure the threshold of each resource group.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
44004	Major	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

If the number of queuing tasks in a resource group exceeds the threshold, a large number of tasks may be in the queuing state. The Presto task time exceeds the expected value. When the number of queuing tasks in a resource group exceeds the maximum number (**maxQueued**) of queuing tasks in the resource group, new tasks cannot be executed.

## Possible Causes

The resource group configuration is improper or too many tasks in the resource group are submitted.

## Handling Procedure

- Step 1** Choose **Components > Presto > Service Configuration** (switch **Basic** to **All**) > **Presto > resource-groups** to adjust the resource group configuration.
- Step 2** You can choose **Components > Presto > Service Configuration** (switch **Basic** to **All**) > **Coordinator > Customize > resourceGroupAlarm** to modify the threshold of each resource group.
- Step 3** Collect the fault information.
  1. Log in to the cluster node based on the host name in the fault information and query the number of queuing tasks based on **Resource Group** in the additional information on the Presto client.

2. Log in to the cluster node based on the host name in the fault information, view the `/var/log/Bigdata/nodeagent/monitorlog/monitor.log` file, and search for resource group information to view the monitoring collection information of the resource group.
3. Contact O&M personnel and send the collected logs.

----End

## Related Information

None

## 7.12.274 ALM-44005 Presto Coordinator Process GC Time Exceeds the Threshold

### Description

The system collects GC time of the Presto Coordinator process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). You can change the threshold by choosing **System > Configure Alarm Threshold > Service > Presto > Coordinator > Presto Process Garbage Collection Time > Garbage Collection Time of the Coordinator Process** on MRS Manager. This alarm is cleared when the Coordinator process GC time is less than or equal to the threshold.

### Attribute

Alarm ID	Alarm Severity	Auto Clear
44005	Major	Yes

### Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

### Impact on the System

If the GC time of the Coordinator process is too long, the Coordinator process running performance will be affected and the Coordinator process will even be unavailable.

## Possible Causes

The heap memory of the Coordinator process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Procedure

### Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.

#### NOTE

For MRS 1.8.10 or earlier, log in to MRS Manager and choose **Alarms**.

2. Select the alarm whose **Alarm ID** is **44005** and then check the role name in **Location** and confirm the IP address of the instance.
3. Choose **Components > Presto > Instances > Coordinator** (business IP address of the instance for which the alarm is generated) > **Customize > Presto Garbage Collection Time**. Click **OK** to view the GC time.
4. Check whether the GC time of the Coordinator process is longer than 5 seconds.
  - If yes, go to **Step 1.5**.
  - If no, go to **Step 2**.
5. Choose **Components > Presto > Service Configuration**, and switch **Basic** to **All**. Choose **Presto > Coordinator**. Increase the value of **-Xmx** (maximum heap memory) in the **JAVA\_OPTS** parameter based on the site requirements.
6. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 2**.

### Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M personnel and send the collected logs.

----End

## Reference

None

## 7.12.275 ALM-44006 Presto Worker Process GC Time Exceeds the Threshold

### Description

The system collects GC time of the Presto Worker process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). You can change the threshold by choosing **System > Configure Alarm Threshold > Service > Presto > Worker > Presto Garbage Collection Time > Garbage Collection Time of the Worker Process** on MRS

Manager. This alarm is cleared when the Worker process GC time is shorter than or equal to the threshold.

## Attribute

Alarm ID	Alarm Severity	Auto Clear
44006	Major	Yes

## Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

## Impact on the System

If the GC time of the Worker process is too long, the Worker process running performance will be affected and the Worker process will even be unavailable.

## Possible Causes

The heap memory of the Worker process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Procedure

### Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.

#### NOTE

For MRS 1.8.10 or earlier, log in to MRS Manager and choose **Alarms**.

2. Select the alarm whose **Alarm ID** is **44006**. Then check the role name in **Location** and confirm the IP address of the instance.
3. Choose **Components** > **Presto** > **Instances** > **Worker** (business IP address of the instance for which the alarm is generated) > **Customize** > **Presto Garbage Collection Time**. Click **OK** to view the GC time.
4. Check whether the GC time of the Worker process is longer than 5 seconds.
  - If yes, go to **Step 1.5**.
  - If no, go to **Step 2**.
5. Choose **Components** > **Presto** > **Service Configuration**, and switch **Basic** to **All**, and choose **Presto** > **Worker** Increase the value of **-Xmx** (maximum heap memory) in the **JAVA\_OPTS** parameter based on the site requirements.

6. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M personnel and send the collected logs.

----End

## Reference

None

## 7.12.276 ALM-45000 HetuEngine Service Unavailable

### Alarm Description

The system checks the HetuEngine service status every 300 seconds. This alarm is generated when the HetuEngine service is unavailable.

This alarm is cleared when the HetuEngine service recovers.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45000	Critical	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System

FusionInsight Manager cannot be used to perform operations on the HetuEngine cluster, and HetuEngine functions are unavailable.



## Possible Causes

- The KrbServer service is abnormal.
- The ZooKeeper service is abnormal.
- The HDFS service is abnormal.
- The Yarn service is abnormal.
- The DBService service is abnormal.
- The Hive service is abnormal.
- There are no HSBroker instances in HetuEngine.

## Handling Procedure

### Check the KrbServer service status.

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarm**.

**Step 2** In the alarm list, check whether the "ALM-25500 KrbServer Service Unavailable" alarm is generated.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

**Step 3** Clear "ALM-25500 KrbServer Service Unavailable" according to the alarm help.


**Step 4** In the alarm list, check whether the alarm "ALM-45000 HetuEngine Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

### Check the ZooKeeper service status.

**Step 5** In the alarm list, check whether the alarm "ALM-12007 Process Fault" is generated.

- If yes, go to [Step 6](#).
- If no, go to [Step 9](#).

**Step 6** In the alarm list, click  in the row that contains the "Process Fault" alarm. Check whether the name of the service for which the alarm is generated is ZooKeeper in **Location Information**.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

**Step 7** Clear "ALM-12007 Process Fault" according to the alarm help.

**Step 8** In the alarm list, check whether the alarm "ALM-45000 HetuEngine Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

### Check the HDFS service status.

**Step 9** In the alarm list, check whether the "ALM-14000 HDFS Service Unavailable" alarm is generated.

- If yes, go to [Step 10](#).
- If no, go to [Step 12](#).

**Step 10** Clear "ALM-14000 HDFS Service Unavailable" according to the alarm help.

**Step 11** In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

**Check the YARN service status.**

**Step 12** In the alarm list, check whether the "ALM-18000 YARN Service Unavailable" alarm is generated.

- If yes, go to [Step 13](#).
- If no, go to [Step 15](#).

**Step 13** Clear "ALM-18000 YARN Service Unavailable" according to the alarm help.

**Step 14** In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 15](#).

**Check the DBService service status.**

**Step 15** In the alarm list, check whether the "ALM-27001 DBService Service Unavailable" alarm is generated.

- If yes, go to [Step 16](#).
- If no, go to [20](#).

**Step 16** Clear "ALM-27001 DBService Service Unavailable" according to the alarm help.

**Step 17** In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to [20](#).

**Check the Hive service status.**

**Step 18** In the alarm list, check whether the "ALM-16004 Hive Service Unavailable" alarm is generated.


- If yes, go to [Step 19](#).
- If no, go to [20](#).

**Step 19** Clear "ALM-16004 Hive Service Unavailable" according to the alarm help.

**Step 20** In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to [20](#).

**Check whether there are no HSBroker instances in HetuEngine.**

- Step 21** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HetuEngine**. On the page that is displayed, click the **Instance** tab.
- Step 22** Check whether there are no HSBroker instances.
- If yes, click **Add Instance** to add one.
  - If no, go to [23](#).
- Step 23** In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.
- If yes, no further action is required.
  - If no, go to [23](#).
- Check the network connection between HetuEngine and ZooKeeper, HDFS, YARN, DBService, and Hive.**
- Step 24** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HetuEngine**. On the page that is displayed, click the **Instance** tab.
- Step 25** Click the host name in the **HSBroker** row and record the management IP address in the **Basic Information** area.
- Step 26** Log in to the host where HSBroker resides as user **omm** using the IP address obtained in [Step 25](#).
- Step 27** Run the **ping** command to check whether the network connection between the host where HSBroker resides and the hosts where ZooKeeper, HDFS, Yarn, DBService, and Hive reside is in the normal state.
- If yes, go to [Step 30](#).
  - If no, go to [Step 28](#).
- Step 28** Contact the network administrator to restore the network.
- Step 29** In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 30](#).
- Collect fault information.**
- Step 30** On FusionInsight Manager, choose **O&M** > **Log** > **Download**.
- Step 31** Expand the **Service** drop-down list. In the **Services** dialog box that is displayed, select **HetuEngine** under the target cluster name, and click **OK**.
- Step 32** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.
- Step 33** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 34** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

After the fault is rectified, the system automatically clears this alarm.

## Reference

None

## 7.12.277 ALM-45001 Faulty HetuEngine Compute Instances

This alarm applies only to MRS 3.2.0 or later.

## Alarm Description

The system checks the HetuEngine compute instance status every 60 seconds. This alarm is generated when a HetuEngine compute instance is faulty.

This alarm is cleared when all faulty HetuEngine compute instances are restored.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45001	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

SQL tasks submitted to the faulty compute instance of HetuEngine fail to be executed.

## Possible Causes

- The HDFS service is abnormal.
- The Yarn service is abnormal.

- Yarn queue resources are insufficient.
- The process of compute instances is faulty.

## Handling Procedure

### Check the HDFS service status.

**Step 1** In the alarm list, check whether the "ALM-14000 HDFS Service Unavailable" alarm is generated.

- If yes, go to [Step 2](#).
- If no, go to [Step 4](#).

**Step 2** Clear "ALM-14000 HDFS Service Unavailable" according to the alarm help.

**Step 3** In the alarm list, check whether the "ALM-45001 Faulty HetuEngine Compute Instances" alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

### Check the YARN service status.

**Step 4** In the alarm list, check whether the "ALM-18000 YARN Service Unavailable" alarm is generated.

- If yes, go to [Step 5](#).
- If no, go to [Step 7](#).

**Step 5** Clear "ALM-18000 YARN Service Unavailable" according to the alarm help.

**Step 6** In the alarm list, check whether the "ALM-45001 Faulty HetuEngine Compute Instances" alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

### Check the YARN queue resource status.

**Step 7** In the alarm list, check whether the "ALM-18022 Insufficient YARN Queue Resources" alarm is generated.

- If yes, go to [8](#).
- If no, go to [Step 10](#).


**Step 8** Clear "ALM-18022 Insufficient YARN Queue Resources" according to the alarm help.

**Step 9** In the alarm list, check whether the "ALM-45001 Faulty HetuEngine Compute Instances" alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

### Check the HetuEngine compute instance status.

**Step 10** Log in to FusionInsight Manager as an administrator who can access the HetuEngine web UI and choose **Cluster > Services > HetuEngine**.

- Step 11** In the **Basic Information** area on the **Dashboard** tab page, click the link next to **HSConsole WebUI** to access the HSConsole page.
- Step 12** On the compute instance page, check whether any compute instances are in the **FAULT** state.
- If yes, go to **Step 13**.
  - If no, go to **Step 14**.
- Step 13** In the **Operation** column of the target compute instance, click **Start** and wait until the instance is started.
- Step 14** In the alarm list, check whether the "ALM-45001 Faulty HetuEngine Compute Instances" alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 15**.
- Collect fault information.**
- Step 15** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 16** Expand the **Service** drop-down list. In the **Services** dialog box that is displayed, select **HetuEngine** under the target cluster name, and click **OK**.
- Step 17** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.
- Step 18** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 19** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.278 ALM-45003 HetuEngine QAS Disk Capacity Is Insufficient

This section applies to MRS 3.3.0 or later.

## Alarm Description

The system checks the HetuEngine QAS disk usage every 60 seconds and compares the actual disk usage with the threshold. The disk usage has a default threshold. This alarm is generated if the disk usage exceeds the threshold.

To change the threshold, choose **O&M > Alarm > Thresholds**. In the service list, choose **HetuEngine > Disk > QAS Disk Usage (QAS)**.

If the **Trigger Count** is **1**, this alarm is cleared when the usage of the HetuEngine QAS disk is less than or equal to the threshold. If the **Trigger Count** is greater than **1**, this alarm is cleared when the disk usage is less than or equal to 80% of the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45003	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Data cannot be written to the HetuEngine QAS disk. SQL diagnosis and materialized view recommendation of HetuEngine SQL O&M are unavailable.

## Possible Causes

- The alarm threshold is improperly configured.
- The configuration of the HetuEngine QAS disk cannot meet service requirements. The disk usage reaches the upper limit.

## Handling Procedure

**Check whether the threshold is set properly.**

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. In the service list, choose **HetuEngine > Disk > QAS Disk Usage (QAS)**. Check whether the alarm threshold is set properly. The default threshold is 80% of the disk capacity. You can change the threshold as required.

- If the threshold is set properly, go to [Step 4](#).
- If the threshold is not set properly, go to [Step 2](#).

**Step 2** Click **Modify** in the **Operation** column to modify and save the alarm threshold as required.

**Step 3** Wait 2 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to [Step 4](#).

**Check whether the disk usage reaches the upper limit.**

**Step 4** Expand the alarm information, view the information in the **Location** area, and check the role name and host name of the QAS disk where the alarm is generated.

**Step 5** Choose **Cluster > Services > HetuEngine** and click **Instance**. On the displayed page, click the QAS role name in the alarm information. On the instance page that is displayed, click **Chart** and check whether the QAS disk usage in the **QAS Disk Usage** chart exceeds the threshold (80% of the disk capacity by default).

- If the disk usage reaches the upper limit, go to [Step 6](#).
- If the disk usage does not reaches the upper limit, go to [Step 9](#).

**Step 6** Log in to the host of the node where the QAS instance reporting the alarm is located as the **root** user.

**Step 7** Run the following command to go to the QAS data directory and delete temporary files as required:

```
cd ${BIGDATA_DATA_HOME}/hetuengine/qas
```

---

**NOTICE**

Deleting temporary files affects the latest QAS execution result but does not affect subsequent results.

---

**Step 8** Wait 2 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to [Step 9](#).

**Collect fault information.**

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Expand the **Service** drop-down list, select **HetuEngine** for the target cluster, and click **OK**.

**Step 11** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

**Step 12** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.



**Step 13** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

# 7.12.279 ALM-45175 Average Time for Calling OBS Metadata APIs Is Greater than the Threshold

## Alarm Description

The system checks whether the average duration for calling OBS metadata APIs is greater than the threshold every 30 seconds. This alarm is generated when the number of consecutive times that the average time exceeds the specified threshold is greater than the number of smoothing times.

This alarm is automatically cleared when the average duration for calling the OBS metadata APIs is lower than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45175	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the average time for calling the OBS metadata APIs exceeds the threshold, the upper-layer big data computing services may be affected. To be more specific, the execution time of some computing tasks will exceed the threshold.

## Possible Causes

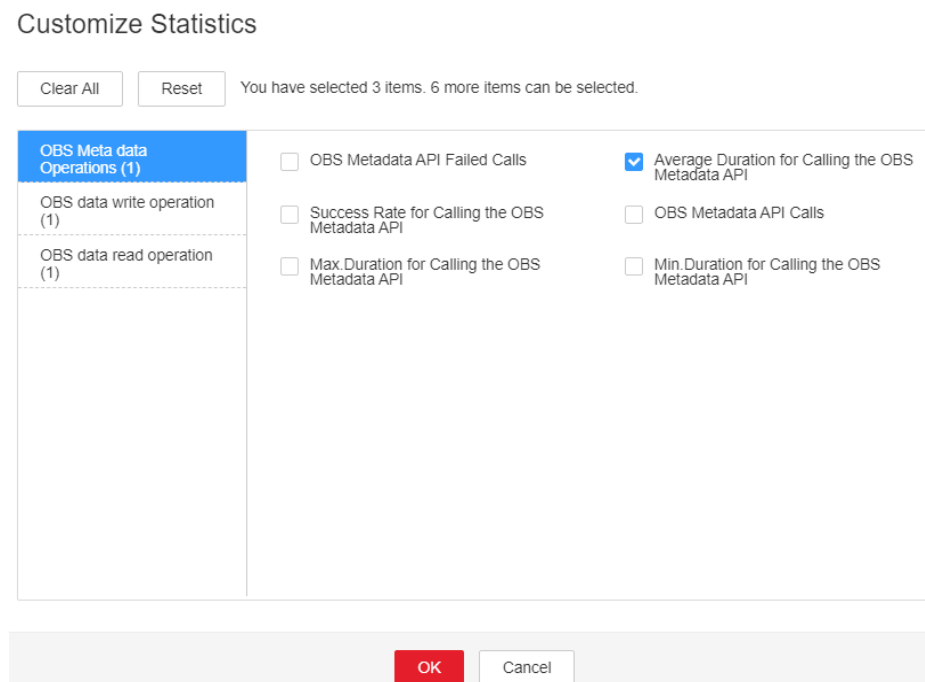
Frame freezing occurs on the OBS server, or the network between the OBS client and the OBS server is unstable.

## Handling Procedure

**Check the heap memory usage.**

- Step 1** On the **FusionInsight Manager** homepage, choose **O&M > Alarm > Alarms > Average Time for Calling the OBS Metadata API Exceeds the Threshold**, view the role name in **Location**, and check the instance IP address.
- Step 2** Choose **Cluster > Name of the desired cluster > Services > meta > Instance > meta** (IP address of the instance for which the alarm is generated). Click the drop-down list in the upper right corner of the chart area and choose **Customize**. In the dialog box that is displayed, select **Average time of OBS interface calls** from **OBS Meta data Operations**, and click **OK**. Check whether the average time of OBS metadata API calls exceeds the threshold.
- If yes, go to [Step 3](#).
  - If no, go to [Step 5](#).

**Figure 7-160** Average duration for calling the OBS metaData API



- Step 3** Choose **Cluster > Name of the desired cluster > O&M > Alarm > Thresholds > meta > Average Time for Calling the OBS Metadata API**. Increase the threshold or smoothing times as required.


**Step 4** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Collect the fault information.**

**Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** In the **Services** area, select **NodeAgent**, **NodeMetricAgent**, **OmmServer**, and **OmmAgent** under OMS.

**Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.280 ALM-45176 Success Rate of Calling OBS Metadata APIs Is Lower than the Threshold

## Alarm Description

The system checks whether the success rate of calling OBS metadata APIs is lower than the threshold every 30 seconds. This alarm is generated when the success rate is lower than the threshold.

This alarm is automatically cleared when the success rate of calling APIs for writing OBS data is greater than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45176	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the success rate of calling the OBS metadata APIs is less than the threshold, the upper-layer big data computing services may be affected. To be more specific, some computing tasks may fail to be executed.

## Possible Causes

An execution exception or severe timeout occurs on the OBS server.

## Handling Procedure

**Check the heap memory usage.**

- Step 1** On the **FusionInsight Manager** homepage, choose **O&M > Alarm > Alarms > Success Rate for Calling the OBS Metadata API Is Lower Than the Threshold**, view the role name in **Location**, and check the instance IP address.
- Step 2** Choose **Cluster > Name of the desired cluster > Services > meta > Instance > meta** (IP address of the instance for which the alarm is generated). Click the drop-down list in the upper right corner of the chart area and choose **Customize**. In the dialog box that is displayed, select **Success percent of OBS interface calls** from **OBS Meta data Operations**, and click **OK**. Check whether the average time of OBS metadata API calls exceeds the threshold.
- If yes, go to [Step 3](#).
  - If no, go to [Step 5](#).

**Figure 7-161** Successful rate for calling the OBS API

Customize Statistics

You have selected 3 items. 6 more items can be selected.

<input checked="" type="checkbox"/> OBS Meta data Operations (1)	<input type="checkbox"/> OBS Metadata API Failed Calls	<input type="checkbox"/> Average Duration for Calling the OBS Metadata API
<input type="checkbox"/> OBS data write operation (1)	<input checked="" type="checkbox"/> Success Rate for Calling the OBS Metadata API	<input type="checkbox"/> OBS Metadata API Calls
<input type="checkbox"/> OBS data read operation (1)	<input type="checkbox"/> Max.Duration for Calling the OBS Metadata API	<input type="checkbox"/> Min.Duration for Calling the OBS Metadata API

**Step 3** Choose **Cluster** > *Name of the desired cluster* > **O&M** > **Alarm** > **Thresholds** > **meta** > **Success Rate for Calling the OBS Metadata API**. Increase the threshold or smoothing times as required.


**Step 4** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Collect the fault information.**

**Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 6** In the **Services** area, select **NodeAgent**, **NodeMetricAgent**, **OmmServer**, and **OmmAgent** under OMS.

**Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

### 7.12.281 ALM-45177 Success Rate of Calling OBS Data Read APIs Is Lower than the Threshold

#### Alarm Description

The system checks whether the success rate of calling APIs for reading OBS data is lower than the threshold every 30 seconds. This alarm is generated when the success rate is lower than the threshold.

This alarm is automatically cleared when the success rate of calling APIs for reading OBS data is greater than the threshold.

#### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45177	Minor	Yes

#### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

#### Impact on the System

If the success rate of calling the OBS APIs for reading data is less than the threshold, the upper-layer big data computing services may be affected. To be more specific, some computing tasks may fail to be executed.

#### Possible Causes

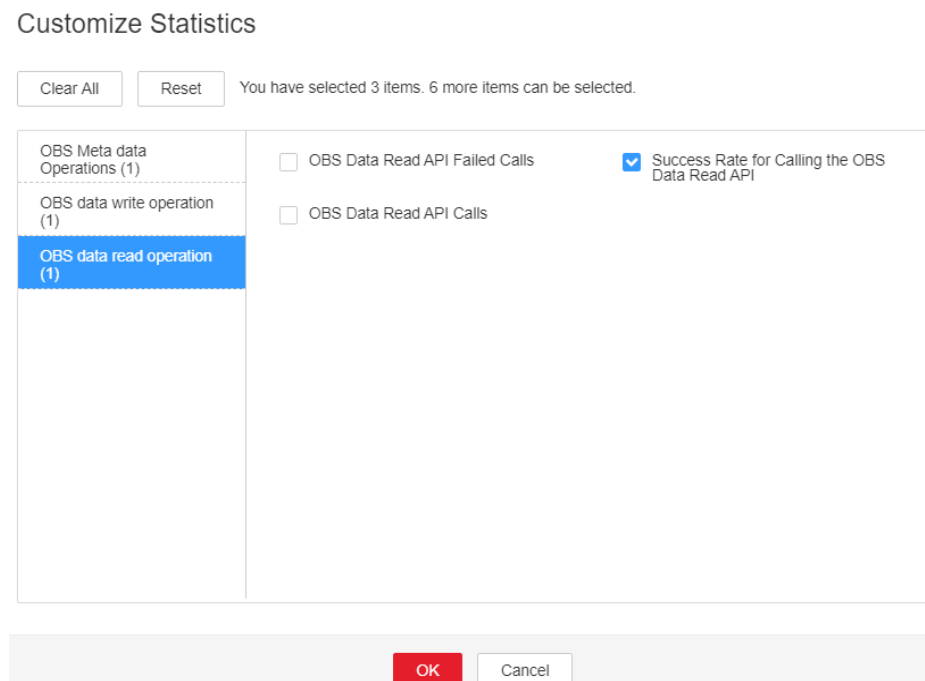
An execution exception or severe timeout occurs on the OBS server.

## Handling Procedure


Check the heap memory usage.

- Step 1** On the **FusionInsight Manager** homepage, choose **O&M > Alarm > Alarms > Success Rate for Calling the OBS Data Read API Is Lower Than the Threshold**, view the role name in **Location**, and check the instance IP address.
- Step 2** Choose **Cluster > Name of the desired cluster > Services > meta > Instance > meta** (IP address of the instance for which the alarm is generated). Click the drop-down list in the upper right corner of the chart area and choose **Customize**. In the dialog box that is displayed, select **Success percent of OBS data read operation interface calls** from **OBS data read operation**, and click **OK**. Check whether the average time of OBS metadata API calls exceeds the threshold.
- If yes, go to [Step 3](#).
  - If no, go to [Step 5](#).

**Figure 7-162** Success rate for calling the OBS data read API



- Step 3** Choose **Cluster > Name of the desired cluster > O&M > Alarm > Thresholds > meta > Success Rate for Calling the OBS Data Read API**. Increase the threshold or smoothing times as required.




- Step 4** Check whether the alarm is cleared.
- If yes, no further action is required.

- If no, go to [Step 5](#).

**Collect the fault information.**

**Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** In the **Services** area, select **NodeAgent**, **NodeMetricAgent**, **OmmServer**, and **OmmAgent** under OMS.

**Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.282 ALM-45178 Success Rate of Calling OBS Data Write APIs Is Lower Than the Threshold

## Alarm Description

The system checks whether the success rate of calling APIs for writing OBS data is lower than the threshold every 30 seconds. This alarm is generated when the success rate is lower than the threshold.

This alarm is automatically cleared when the success rate of calling APIs for writing OBS data is greater than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45178	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.



Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the success rate of calling the OBS APIs for writing data is lower than the threshold, the upper-layer big data computing services may be affected. To be more specific, some computing tasks may fail to be executed.

## Possible Causes

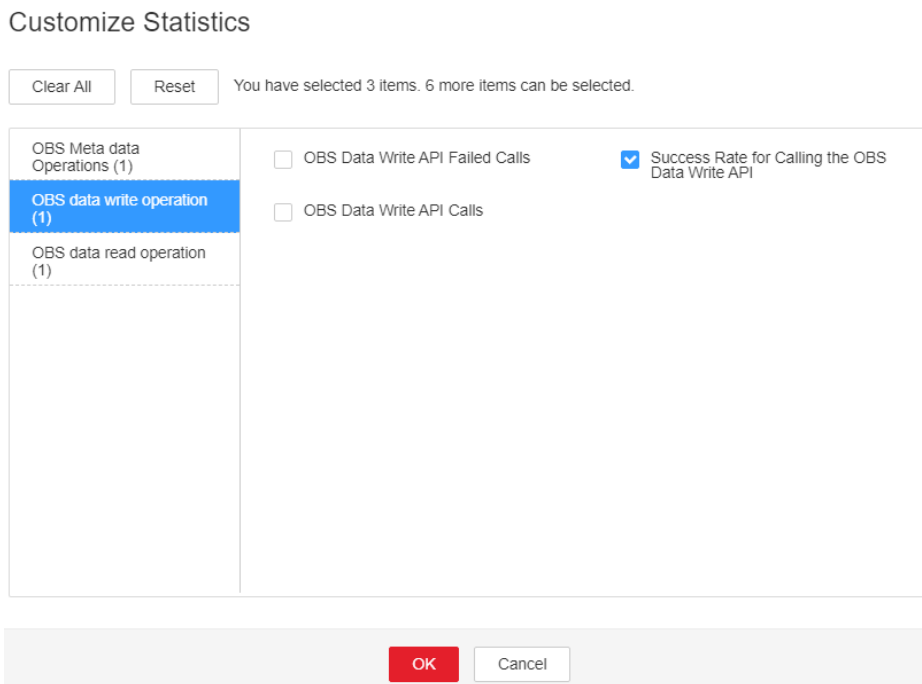
An execution exception or severe timeout occurs on the OBS server.

## Handling Procedure


**Check the heap memory usage.**

- Step 1** On the **FusionInsight Manager** homepage, choose **O&M > Alarm > Alarms > Success Rate for Calling the OBS Data Write API Is Lower Than the Threshold**, view the role name in **Location**, and check the instance IP address.
- Step 2** Choose **Cluster > Name of the desired cluster > Services > meta > Instance > meta** (IP address of the instance for which the alarm is generated). Click the drop-down list in the upper right corner of the chart area and choose **Customize**. In the dialog box that is displayed, select **Success percent of OBS data write operation interface calls** from **OBS data write operation**, and click **OK**. Check whether the average time of OBS metadata API calls exceeds the threshold.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 5](#).

**Figure 7-163** Success rate for calling the OBS data write API



**Step 3** Choose **Cluster > Name of the desired cluster > O&M > Alarm > Thresholds > meta > Success Rate for Calling the OBS Data Write API**. Increase the threshold or smoothing times as required.




**Step 4** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Collect the fault information.**

**Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 6** In the **Services** area, select **NodeAgent**, **NodeMetricAgent**, **OmmServer**, and **OmmAgent** under OMS.

**Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.283 ALM-45179 Number of Failed OBS readFully API Calls Exceeds the Threshold

## Alarm Description

The system checks whether the number of failed OBS readFully API calls exceeds the threshold every 30 seconds. This alarm is generated when the number of failed API calls exceeds the threshold.

This alarm is automatically cleared when the number of failed OBS readFully API calls is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45179	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Certain upper-layer big data computing tasks will fail to execute.

## Possible Causes


An execution exception or severe timeout occurs on the OBS server.

## Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **meta > Number of failed calls to the OBS readFully interface**. In the right pane, set **Threshold** or **Trigger Count** to a larger value as required.
- Step 2** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 3**.
- Step 3** Contact OBS O&M personnel to check whether the OBS service is normal.
- If yes, go to **Step 4**.
  - If no, contact OBS O&M personnel to restore the OBS service.


### Collect fault information.

- Step 4** On FusionInsight Manager, choose **Cluster > Services > meta**. On the page that is displayed, click the **Chart** tab. On this tab page, select **OBS data read operation** in the **Chart Category** area. In the **Number of failed calls to the OBS readFully interface-All Instances** chart, view the host name of the instance that has the maximum number of failed OBS readFully API calls. For example, the host name is **node-ana-corevpeO003**.



- Step 5** Choose **O&M > Log > Download** and select **meta** and **meta** under it for **Service**.

- Step 6** Select the host obtained in **Step 4** for **Hosts**.

- Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

- Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.284 ALM-45180 Number of Failed OBS read API Calls Exceeds the Threshold

## Alarm Description

The system checks whether the number of failed OBS read API calls exceeds the threshold every 30 seconds. This alarm is generated when the number of failed API calls exceeds the threshold.

This alarm is automatically cleared when the number of failed OBS read API calls is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45180	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Certain upper-layer big data computing tasks will fail to execute.

## Possible Causes

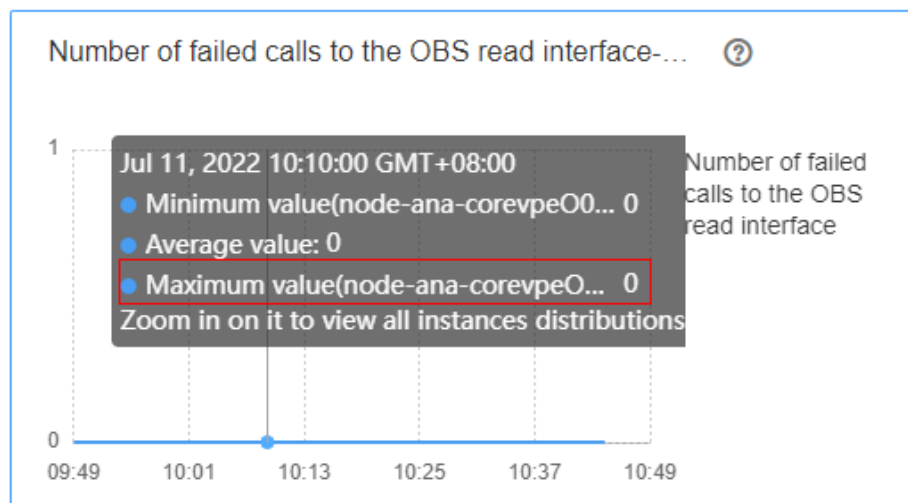
An execution exception or severe timeout occurs on the OBS server.


## Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **meta > Number of failed calls to the OBS read interface**. In the right pane, set **Threshold** or **Trigger Count** to a larger value as required.
- Step 2** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 3**.
- Step 3** Contact OBS O&M personnel to check whether the OBS service is normal.
- If yes, go to **Step 4**.
  - If no, contact OBS O&M personnel to restore the OBS service.

### Collect fault information.

- Step 4** On FusionInsight Manager, choose **Cluster > Services > meta**. On the page that is displayed, click the **Chart** tab. On this tab page, select **OBS data read operation** in the **Chart Category** area. In the **Number of failed calls to the OBS read interface-All Instances** chart, view the host name of the instance that has the maximum number of failed OBS read API calls. For example, the host name is **node-ana-corevpeO003**.



- Step 5** Choose **O&M > Log > Download** and select **meta** and **meta** under it for **Service**.
- Step 6** Select the host obtained in **Step 4** for **Hosts**.
- Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.285 ALM-45181 Number of Failed OBS write API Calls Exceeds the Threshold

## Alarm Description

The system checks whether the number of failed OBS write API calls exceeds the threshold every 30 seconds. This alarm is generated when the number of failed API calls exceeds the threshold.

This alarm is automatically cleared when the number of failed OBS write API calls is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45181	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Certain upper-layer big data computing tasks will fail to execute.

## Possible Causes

An execution exception or severe timeout occurs on the OBS server.

## Handling Procedure

**Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **meta > Number of failed calls to the OBS write interface**. In the right pane, set **Threshold** or **Trigger Count** to a larger value as required.

**Step 2** Check whether the alarm is cleared.

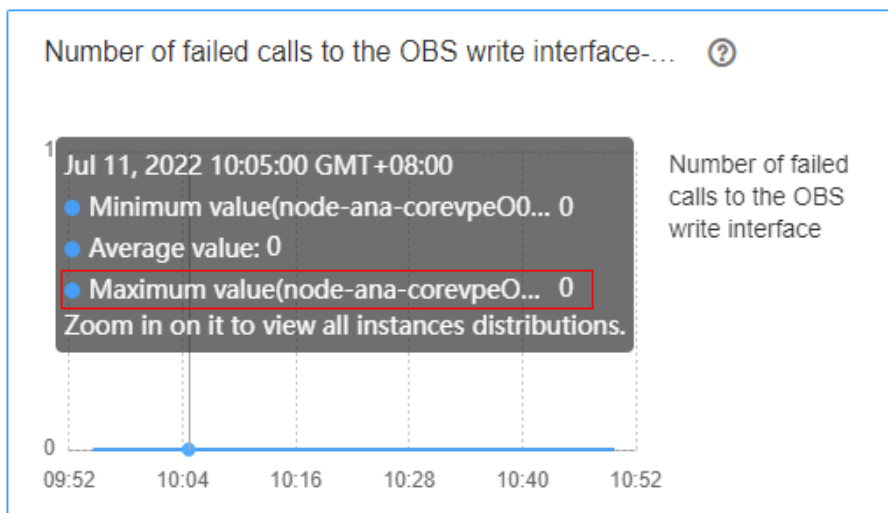
- If yes, no further action is required.
- If no, go to [Step 3](#).

**Step 3** Contact OBS O&M personnel to check whether the OBS service is normal.

- If yes, go to [Step 4](#).
- If no, contact OBS O&M personnel to restore the OBS service.


### Collect fault information.

**Step 4** On FusionInsight Manager, choose **Cluster > Services > meta**. On the page that is displayed, click the **Chart** tab. On this tab page, select **OBS data write operation** in the **Chart Category** area. In the **Number of failed calls to the OBS write interface-All Instances** chart, view the host name of the instance that has the maximum number of failed OBS write API calls. For example, the host name is **node-ana-corevpeO003**.



**Step 5** Choose **O&M > Log > Download** and select **meta** and **meta** under it for **Service**.

**Step 6** Select the host obtained in [Step 4](#) for **Hosts**.

**Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.



**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.286 ALM-45182 Number of Throttled OBS Operations Exceeds the Threshold

## Alarm Description

The system checks whether the number of throttled OBS operations exceeds the threshold every 30 seconds. This alarm is generated when the number of throttled OBS operations exceeds the threshold.

This alarm is automatically cleared when the number of throttled OBS operations is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45182	Minor	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Certain upper-layer big data computing tasks will fail to execute.

## Possible Causes

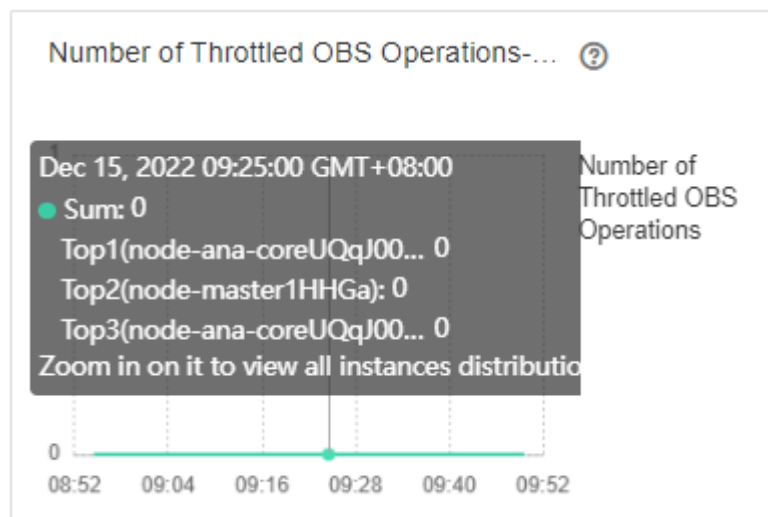
The frequency of requesting OBS APIs is too high.


## Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **meta > Number of Throttled OBS Operations**. In the right pane, set **Threshold** or **Trigger Count** to a larger value as required.
- Step 2** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 3**.
- Step 3** Contact OBS O&M personnel to check whether the OBS service is normal.
- If yes, go to **Step 4**.
  - If no, contact OBS O&M personnel to restore the OBS service.

### Collect fault information.

- Step 4** On FusionInsight Manager, choose **Cluster > Services > meta**. On the page that is displayed, click the **Chart** tab. On this tab page, select **OBS Throttled** in the **Chart Category** area. In the **Number of Throttled OBS Operations-All Instances** chart, view the host name of the instance that has the maximum number of throttled OBS API calls. For example, the host name is **node-ana-coreUQqJ0002**.



- Step 5** Choose **O&M > Log > Download** and select **meta** and **meta** under it for **Service**.
- Step 6** Select the host obtained in **Step 4** for **Hosts**.
- Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.287 ALM-45275 Ranger Service Unavailable

## Alarm Description

The alarm module checks the Ranger service status every 180 seconds. This alarm is generated if the Ranger service is abnormal.

This alarm is cleared after the Ranger service recovers.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45275	Critical	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

When the Ranger service is unavailable, Ranger cannot work properly and the native Ranger UI cannot be accessed.

## Possible Causes

- The DBService service on which Ranger depends is abnormal.
- The RangerAdmin role instance is abnormal.

## Handling Procedure

### Check the DBService process status.

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. On the displayed page, check whether the ALM-27001 DBService Service Unavailable alarm is reported.

- If yes, go to [Step 2](#).
- If no, go to [Step 3](#).

**Step 2** Rectify the DBService service fault by following the handling procedure of ALM-27001 DBService Service Unavailable. After the DBService alarm is cleared, check whether Ranger Service Unavailable alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

### Check all RangerAdmin instances.

**Step 3** Log in to the node where the RangerAdmin instance is located as user **omm** and run the **ps -ef|grep "proc\_rangeradmin"** command to check whether the RangerAdmin process exists on the current node.

- If yes, go to [Step 5](#).
- If no, restart the faulty RangerAdmin instance or Ranger service and go to [Step 4](#).

---

### NOTICE

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---


**Step 4** In the alarm list, check whether the alarm "Ranger Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

### Collect the fault information.

**Step 5** On FusionInsight Manager, choose **O&M > Log > Download**.

**Step 6** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

**Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 8** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

## Related Information

None

# 7.12.288 ALM-45276 Abnormal RangerAdmin Status

## Alarm Description

The alarm module checks the RangerAdmin service status every 60 seconds. This alarm is generated if RangerAdmin is unavailable.

This alarm is automatically cleared after the RangerAdmin service recovers.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45276	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Role for which the alarm is generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System


If the status of a RangerAdmin is abnormal, access to the Ranger native UI is not affected. If there are two abnormal RangerAdmin instances, the Ranger native UI cannot be accessed and operations such as creating, modifying, and deleting policies are unavailable.

## Possible Causes

The RangerAdmin port is not started.

## Handling Procedure

### Check the port process.

- Step 1** In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click  to view the name of the host for which the alarm is generated.
- Step 2** Log in to the node where the RangerAdmin instance is located as user **omm**. Run the `ps -ef|grep "proc_rangeradmin" | grep -v grep | awk -F ' ' '{print $2}'` command to obtain *pid* of the RangerAdmin process, and run the `netstat -anp| grep pid | grep LISTEN` command to check whether the RangerAdmin process monitors port 21401 in a security cluster or port 21400 in a normal cluster.
- If yes, go to [Step 4](#).
  - If no, restart the faulty RangerAdmin instance or Ranger service and go to [Step 3](#).

---


### NOTICE

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---

- Step 3** In the alarm list, check whether the "Abnormal RangerAdmin Status" alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 4](#).

### Collect the fault information.

- Step 4** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 5** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 7** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

## Related Information

None

## 7.12.289 ALM-45277 RangerAdmin Heap Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the heap memory usage of the RangerAdmin service every 60 seconds. This alarm is generated when the system detects that the heap memory usage of the RangerAdmin instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45277	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

Heap memory overflow may cause service breakdown.

### Possible Causes

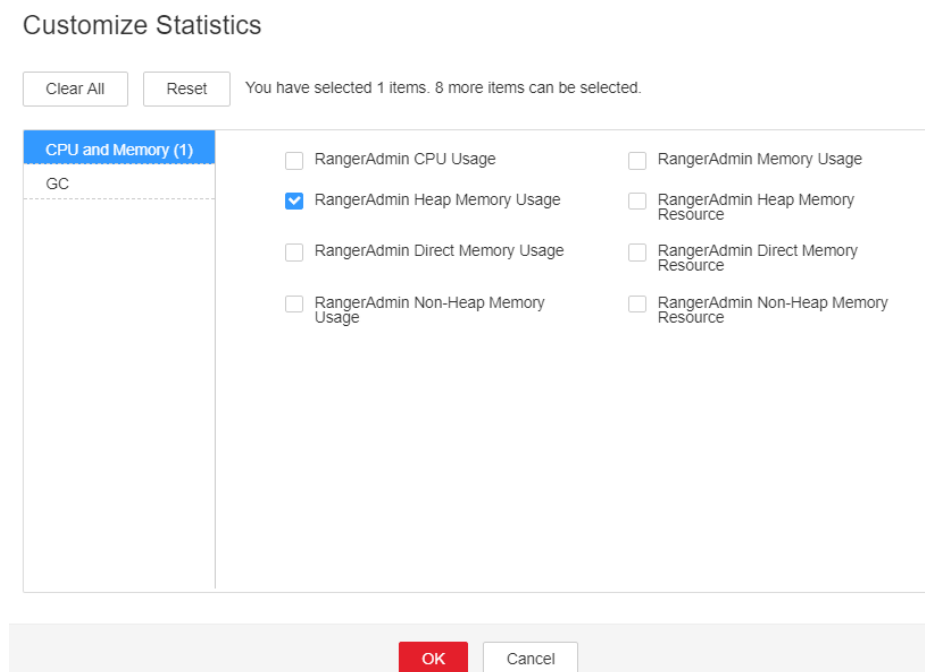
The heap memory usage of the RangerAdmin instance is high or the heap memory is improperly allocated.

### Handling Procedure

**Check the heap memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45277 RangerAdmin Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > RangerAdmin Heap Memory Usage**. Click **OK**.

**Figure 7-164** RangerAdmin heap memory usage



- Step 3** Check whether the heap memory used by RangerAdmin reaches the threshold (95% of the maximum heap memory by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > RangerAdmin > Instance Configuration**. Click **All Configurations**, and choose **RangerAdmin > System**. Increase the value of **-Xmx** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.

**NOTE**

If this alarm is generated, the heap memory configured for RangerAdmin cannot meet the heap memory required by the RangerAdmin process. You are advised to check the heap memory usage of RangerAdmin and change the value of **-Xmx** in **GC\_OPTS** to the twice of the heap memory used by RangerAdmin. The value can be changed based on the actual service scenario. For details, see **Step 2**.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.



- If no, go to [Step 6](#).


---

**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.290 ALM-45278 RangerAdmin Direct Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the direct memory usage of the RangerAdmin service every 60 seconds. This alarm is generated when the direct memory usage of the RangerAdmin instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the direct memory usage of RangerAdmin is less than or equal to the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45278	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Direct memory overflow may cause service breakdown.

## Possible Causes


The direct memory of the RangerAdmin instance is overused or the direct memory is inappropriately allocated. As a result, the memory usage exceeds the threshold.

## Handling Procedure

**Check the direct memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45278 RangerAdmin Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > RangerAdmin Direct Memory Usage**. Click **OK**.

**Figure 7-165** RangerAdmin direct memory usage



**Step 3** Check whether the direct memory used by RangerAdmin reaches the threshold (80% of the maximum direct memory by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > RangerAdmin > Instance Configuration**. Click **All Configurations**, and choose **RangerAdmin > System**. Increase the value of **-XX:MaxDirectMemorySize** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.

**NOTE**

If this alarm is generated, the direct memory configured for RangerAdmin cannot meet the direct memory required by the RangerAdmin process. You are advised to check the direct memory usage of RangerAdmin and change the value of **-XX:MaxDirectMemorySize** in **GC\_OPTS** to the twice of the direct memory used by RangerAdmin. You can change the value based on the actual service scenario. For details, see [Step 2](#).


**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.291 ALM-45279 RangerAdmin Non Heap Memory Usage Exceeds the Threshold

## Alarm Description

The system checks the non-heap memory usage of the RangerAdmin service every 60 seconds. This alarm is generated when the non-heap memory usage of the RangerAdmin instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45279	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Non-heap memory overflow may cause service breakdown.

## Possible Causes


The non-heap memory usage of the RangerAdmin instance is high or the non-heap memory is improperly allocated.

## Handling Procedure

**Check non-heap memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45279 RangerAdmin Non Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > RangerAdmin Non Heap Memory Usage**. Click **OK**.

**Figure 7-166** RangerAdmin non-heap memory usage



**Step 3** Check whether the non-heap memory used by RangerAdmin reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > RangerAdmin > Instance Configuration**. Click **All Configurations**, and choose **RangerAdmin > System**. Set **-XX:MaxPermSize** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

 **NOTE**

If this alarm is generated, the non-heap memory size configured for the RangerAdmin instance cannot meet the non-heap memory required by the RangerAdmin process. You are advised to change the value of **-XX:MaxPermSize** in **GC\_OPTS** to the twice of the current non-heap memory usage or change the value based on the site requirements.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

---

**NOTICE**


During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---

**Collect the fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.292 ALM-45280 RangerAdmin GC Duration Exceeds the Threshold

### Alarm Description

The system checks the GC duration of the RangerAdmin process every 60 seconds. This alarm is generated when the GC duration of the RangerAdmin process exceeds the threshold (12 seconds by default) for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45280	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

The RangerAdmin responds slowly.

### Possible Causes

The heap memory of the RangerAdmin instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

### Handling Procedure

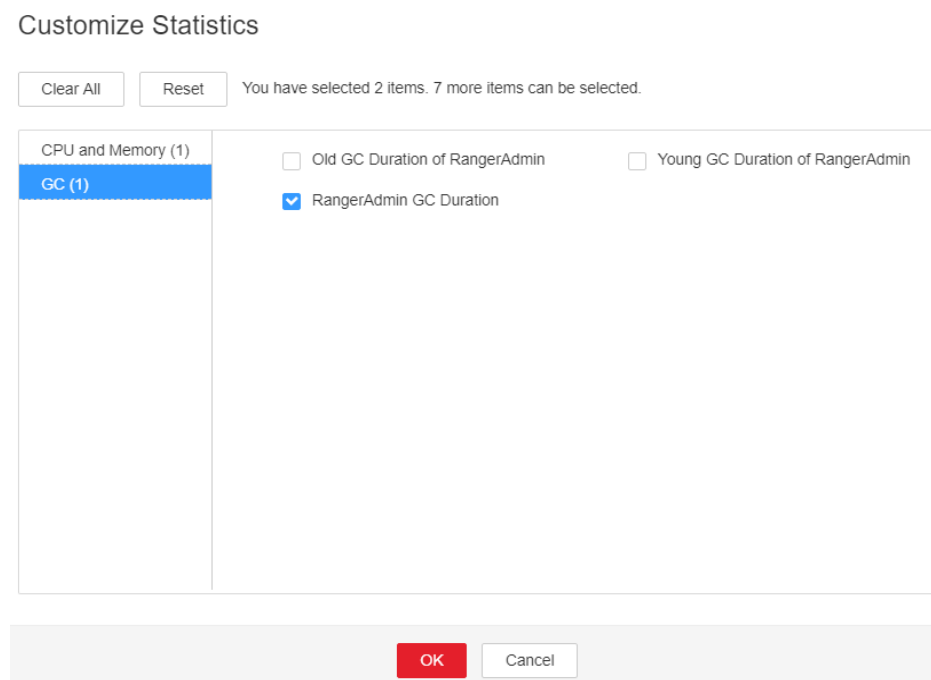
**Check the GC duration.**

**Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45280 RangerAdmin GC Duration Exceeds the Threshold**. Check the location

information of the alarm and view the host name of the instance for which the alarm is generated.

- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize > GC > RangerAdmin GC Duration**. Click **OK**.

**Figure 7-167** RangerAdmin garbage collection (GC) duration



- Step 3** Check whether the GC duration of the RangerAdmin process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

- Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > RangerAdmin > Instance Configuration**. Click **All Configurations**, and choose **RangerAdmin > System**. Increase the value of **-Xmx** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.

**NOTE**

If this alarm is generated, the heap memory configured for RangerAdmin cannot meet the heap memory required by the RangerAdmin process. You are advised to check the heap memory usage of RangerAdmin and change the value of **-Xmx** in **GC\_OPTS** to the twice of the heap memory used by RangerAdmin. The value can be changed based on the actual service scenario. For details, see **Step 2**.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to **Step 6**.



**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

**Alarm Clearance**

This alarm is automatically cleared after the fault is rectified.

**Related Information**

None

## 7.12.293 ALM-45281 UserSync Heap Memory Usage Exceeds the Threshold

**Alarm Description**

The system checks the heap memory usage of the UserSync service every 60 seconds. This alarm is generated when the system detects that the heap memory usage of the UserSync instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

**Alarm Attributes**

Alarm ID	Alarm Severity	Auto Cleared
45281	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

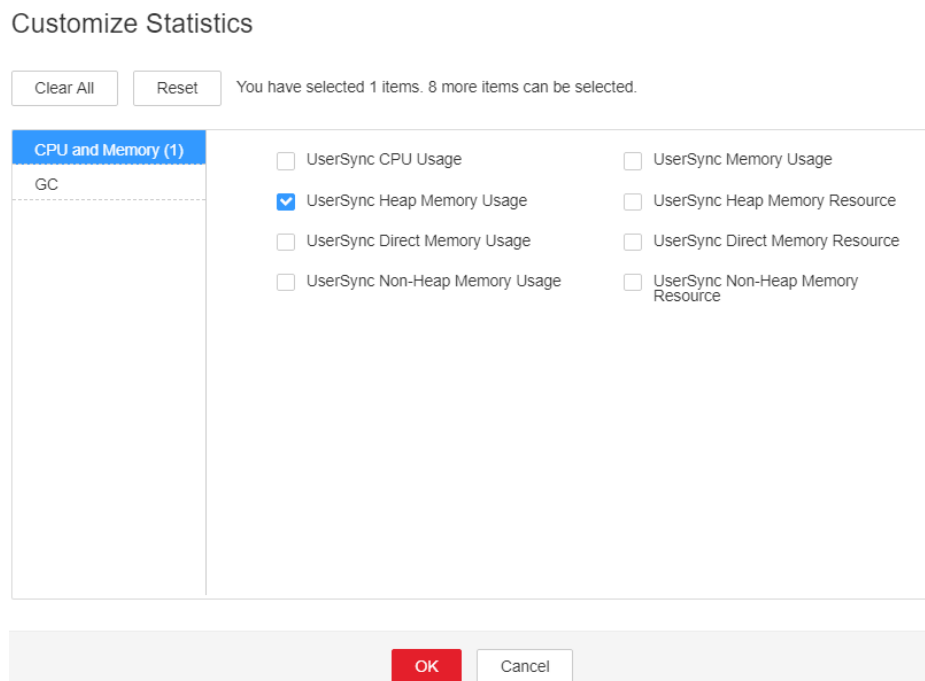
Heap memory overflow may cause service breakdown.

## Possible Causes

The heap memory usage of the UserSync instance is high or the heap memory is improperly allocated.

## Handling Procedure

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45281 UserSync Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > UserSync Heap Memory Usage**. Click **OK**.

**Figure 7-168** UserSync heap memory usage

**Step 3** Check whether the heap memory used by UserSync reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > UserSync > Instance Configuration**. Click **All Configurations**, and choose **UserSync > System**. Increase the value of **-Xmx** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.

**NOTE**

If this alarm is generated, the heap memory configured for UserSync cannot meet the heap memory required by the UserSync process. You are advised to change the **-Xmx** value of **GC\_OPTS** to twice that of the heap memory used by UserSync. You can change the value based on the actual service scenario. For details about how to check the UserSync heap memory usage, see [Step 2](#).

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).


---

**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.294 ALM-45282 UserSync Direct Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the direct memory usage of the UserSync service every 60 seconds. This alarm is generated when the direct memory usage of the UserSync instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the UserSync direct memory usage is less than or equal to the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45282	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Direct memory overflow may cause service breakdown.

## Possible Causes

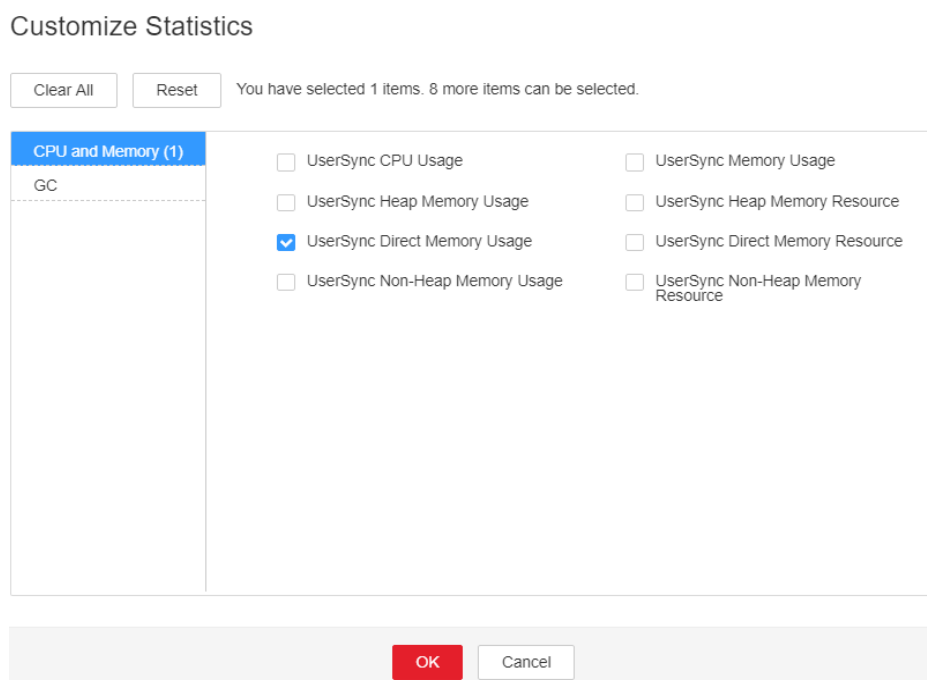
The direct memory of the UserSync instance is overused or the direct memory is inappropriately allocated. As a result, the memory usage exceeds the threshold.


## Handling Procedure

**Check the direct memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45282 UserSync Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm. Check the name of the instance host for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > UserSync Direct Memory Usage**. Click **OK**.

**Figure 7-169** UserSync direct memory usage



- Step 3** Check whether the direct memory used by the UserSync reaches the threshold (80% of the maximum direct memory by default).
- If yes, go to [Step 4](#).
  - If no, go to [Step 6](#).
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > UserSync > Instance Configuration**. Click **All Configurations**, and choose **UserSync > System**. Increase the value of **-XX:MaxDirectMemorySize** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.
-  **NOTE**
- If this alarm is generated, the direct memory configured for UserSync cannot meet the direct memory required by the UserSync process. You are advised to check the direct memory usage of UserSync and change the value of **-XX:MaxDirectMemorySize** in **GC\_OPTS** to the twice of the direct memory used by UserSync. You can change the value based on the actual service scenario. For details, see [Step 2](#).
- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).


---

**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.295 ALM-45283 UserSync Non Heap Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the non-heap memory usage of the UserSync service every 60 seconds. This alarm is generated when the non-heap memory usage of the UserSync instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45283	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

Non-heap memory overflow may cause service breakdown.

### Possible Causes

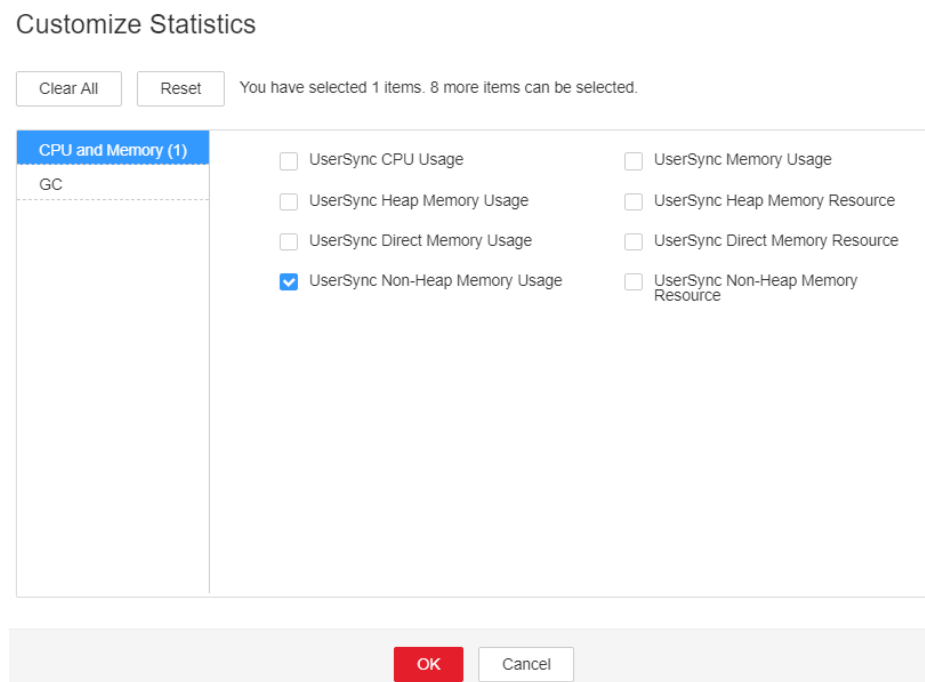
The non-heap memory of the UserSync process is overused or the non-heap memory is inappropriately allocated.

### Handling Procedure

**Check non-heap memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45283 UserSync Non Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > UserSync Non Heap Memory Usage**. Click **OK**.

**Figure 7-170** UserSync non-heap memory usage



- Step 3** Check whether the non-heap memory used by UserSync reaches the threshold (80% of the maximum non-heap memory by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > UserSync > Instance Configuration**. Click **All Configurations**, and choose **UserSync > System**. Set **-XX:MaxPermSize** in the **GC\_OPTS** parameter to a larger value based on site requirements and click **Save** to save the configuration.

**NOTE**

If this alarm is generated, the non-heap memory size configured for the UserSync instance cannot meet the non-heap memory required by the UserSync process. You are advised to change the **-XX:MaxPermSize** value of **GC\_OPTS** to twice that of the current non-heap memory size or change the value based on the site requirements.


- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 6**.



**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

**Alarm Clearance**

This alarm is automatically cleared after the fault is rectified.

**Related Information**

None

**7.12.296 ALM-45284 UserSync Garbage Collection (GC) Time Exceeds the Threshold**

**Alarm Description**

The system checks the GC duration of the UserSync process every 60 seconds. This alarm is generated when the GC duration of the UserSync process exceeds the threshold (12 seconds by default) for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

**Alarm Attributes**

Alarm ID	Alarm Severity	Auto Cleared
45284	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Threshold for triggering the alarm.

## Impact on the System

UserSync responds slowly.

## Possible Causes

The heap memory of the UserSync instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

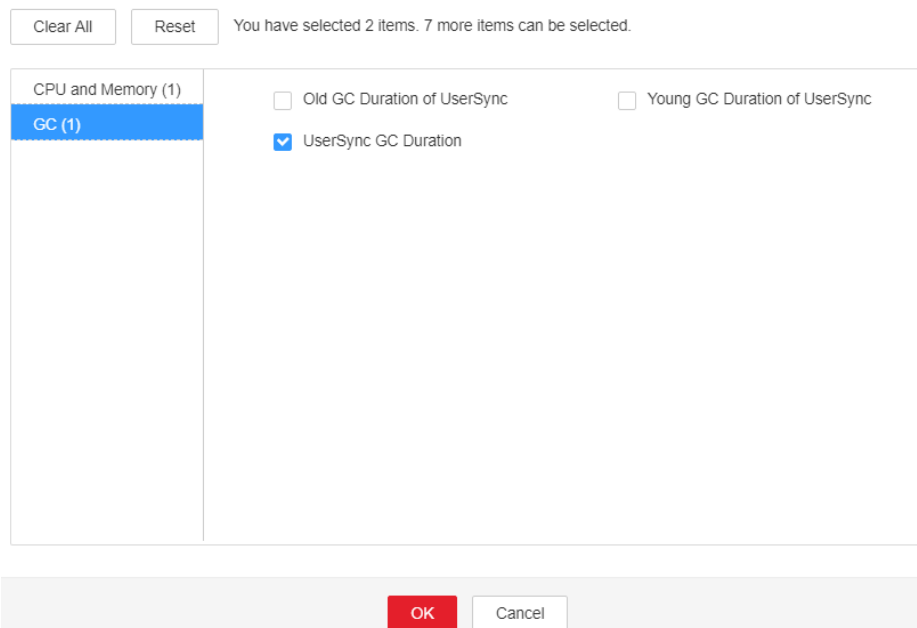
## Handling Procedure

**Check the GC time.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45284 UserSync Garbage Collection (GC) Time Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize > GC > UserSync GC Duration**. Click **OK**.

**Figure 7-171** UserSync GC Duration

Customize Statistics



**Step 3** Check whether the GC duration of the UserSync process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > UserSync > Instance Configuration**. Click **All Configurations**, and choose **UserSync > System**. Increase the value of **-Xmx** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.

**NOTE**

If this alarm is generated, the heap memory configured for UserSync cannot meet the heap memory required by the UserSync process. You are advised to change the value of **GC\_OPTS** to the twice that of the heap memory used by UserSync. You can change the value based on the actual service scenario. For details about how to check the UserSync heap memory usage, see [Step 2](#).


**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

## Related Information

None

## 7.12.297 ALM-45285 TagSync Heap Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the heap memory usage of the TagSync service every 60 seconds. This alarm is generated when the heap memory usage of the TagSync instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45285	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Heap memory overflow may cause service breakdown.

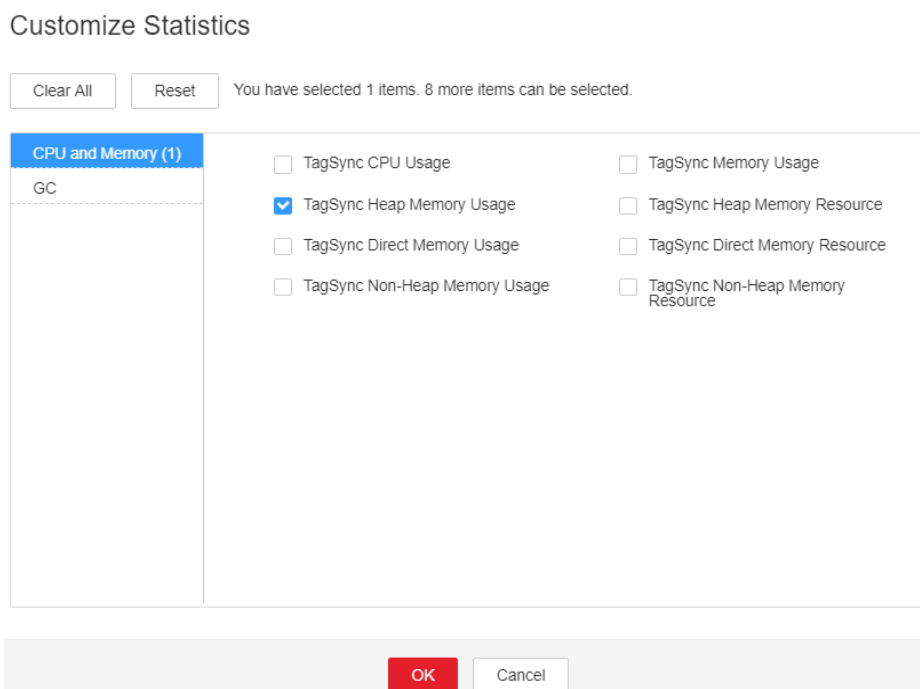
## Possible Causes


The heap memory usage of the TagSync instance is high or the heap memory is improperly allocated.

## Handling Procedure

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45285 TagSync Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > TagSync Heap Memory Usage**. Click **OK**.

**Figure 7-172** TagSync heap memory usage



- Step 3** Check whether the heap memory used by TagSync reaches the threshold (95% of the maximum heap memory by default).
- If yes, go to [Step 4](#).
  - If no, go to [Step 6](#).
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > TagSync > Instance Configuration**. Click **All Configurations** and choose **TagSync > System**. Increase the value of **-Xmx** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.
-  **NOTE**
- If this alarm is generated, the heap memory configured for TagSync cannot meet the heap memory required by the TagSync process. You are advised to change the **-Xmx** value of **GC\_OPTS** to twice that of the heap memory used by TagSync. You can change the value based on the actual service scenario. For details about how to check the TagSync heap memory usage, see [Step 2](#).
- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).


---

**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.298 ALM-45286 TagSync Direct Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the direct memory usage of the TagSync service every 60 seconds. This alarm is generated when the direct memory usage of the TagSync instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the TagSync direct memory usage is less than or equal to the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45286	Major	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

Direct memory overflow may cause service breakdown.

### Possible Causes


The direct memory of the TagSync instance is overused or the direct memory is inappropriately allocated. As a result, the memory usage exceeds the threshold.

### Handling Procedure

**Check the direct memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45286 TagSync Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > TagSync Direct Memory Usage**. Click **OK**.

**Figure 7-173** TagSync direct memory usage



- Step 3** Check whether the direct memory used by the TagSync reaches the threshold (80% of the maximum direct memory by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > TagSync > Instance Configuration**. Click **All Configurations** and choose **TagSync > System**. Increase the value of **-XX:MaxDirectMemorySize** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.

**NOTE**

If this alarm is generated, the direct memory configured for TagSync cannot meet the direct memory required by the TagSync process. You are advised to check the direct memory usage of TagSync and change the value of **-XX:MaxDirectMemorySize** in **GC\_OPTS** to the twice of the direct memory used by TagSync. You can change the value based on the actual service scenario. For details, see **Step 2**.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.



- If no, go to [Step 6](#).


---

**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---

**Collect the fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.299 ALM-45287 TagSync Non Heap Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the non-heap memory usage of the TagSync service every 60 seconds. This alarm is generated when the non-heap memory usage of the TagSync instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45287	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Non-heap memory overflow may cause service breakdown.

## Possible Causes

The non-heap memory of the TagSync process is overused or the non-heap memory is inappropriately allocated.

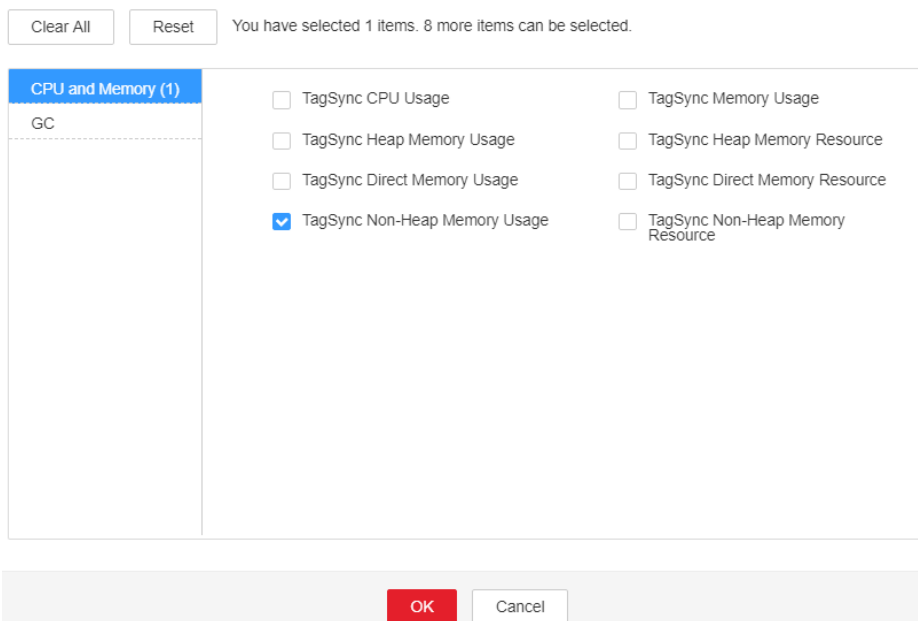
## Handling Procedure

**Check non-heap memory usage.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45287 TagSync Non Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > TagSync Non Heap Memory Usage**. Click **OK**.

**Figure 7-174** TagSync non-heap memory usage

Customize Statistics



**Step 3** Check whether the non-heap memory used by TagSync reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > TagSync > Instance Configuration**. Click **All Configurations** and choose **TagSync > System**. Set **-XX:MaxPermSize** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

**NOTE**

If this alarm is generated, the non-heap memory size configured for the TagSync instance cannot meet the non-heap memory required by the TagSync process. You are advised to change the **-XX:MaxPermSize** value of **GC\_OPTS** to twice that of the current non-heap memory size or change the value based on the site requirements.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).


**NOTICE**

During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

**Collect the fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.300 ALM-45288 TagSync Garbage Collection (GC) Time Exceeds the Threshold

## Alarm Description

The system checks the GC duration of the TagSync process every 60 seconds. This alarm is generated when the GC duration of the TagSync process exceeds the threshold (12 seconds by default) for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45288	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

TagSync responds slowly.

## Possible Causes

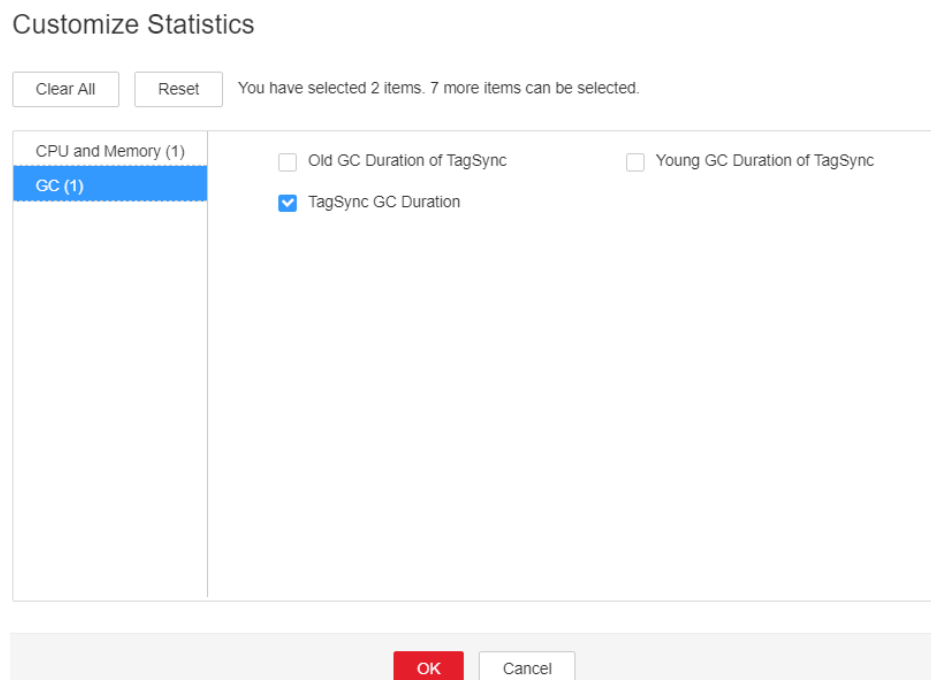
The heap memory of the TagSync instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

## Handling Procedure

**Check the GC duration.**

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45288 TagSync Garbage Collection (GC) Time Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize > GC > TagSync GC Duration**. Click **OK**.

**Figure 7-175** TagSync GC duration



- Step 3** Check whether the GC duration of the TagSync process collected every minute exceeds the threshold (12 seconds by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > TagSync > Instance Configuration**. Click **All Configurations** and choose **TagSync > System**. Increase the value of **-Xmx** in the **GC\_OPTS** parameter based on the site requirements and save the configuration.

 **NOTE**

If this alarm is generated, the heap memory configured for TagSync cannot meet the heap memory required by the TagSync process. You are advised to change the **-Xmx** value of **GC\_OPTS** to twice that of the heap memory used by TagSync. You can change the value based on the actual service scenario. For details about how to check the TagSync heap memory usage, see [Step 2](#).

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

---

**NOTICE**


During the service restart, the service is interrupted. During the instance restart, the instance is unavailable, and the task on that instance fails to be executed.

---

**Collect the fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.301 ALM-45289 PolicySync Heap Memory Usage Exceeds the Threshold

## Alarm Description

The system checks the heap memory usage of the PolicySync service every 60 seconds. This alarm is generated when the heap memory usage of the PolicySync

instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45289	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Heap memory overflow may cause service breakdown.

## Possible Causes

The heap memory of the PolicySync instance is overused or the heap memory is inappropriately allocated.

## Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms > ALM-45289 PolicySync Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > PolicySync Heap Memory Usage**. Click **OK**.

**Figure 7-176** PolicySync Heap Memory Usage

Customize Statistics

Clear All   Reset   You have selected 1 items. 8 more items can be selected.

CPU and Memory(1)		
GC	<input type="checkbox"/> PolicySync CPU Usage	<input type="checkbox"/> PolicySync Memory Usage
	<input checked="" type="checkbox"/> PolicySync Heap Memory Usage	<input type="checkbox"/> PolicySync Heap Memory Resource
	<input type="checkbox"/> PolicySync Direct Memory Usage	<input type="checkbox"/> PolicySync Direct Memory Resource
	<input type="checkbox"/> PolicySync Non-Heap Memory Usage	<input type="checkbox"/> PolicySync Non-Heap Memory Resource

OK   Cancel

**Step 3** Check whether the heap memory used by PolicySync reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > PolicySync**. Click **Instance Configuration** and then **All Configurations**, and choose **PolicySync > System**. Set **-Xmx** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

**NOTE**

If this alarm is generated, the heap memory configured for PolicySync cannot meet the heap memory required by the PolicySync process. You are advised to change the value of **-Xmx** in **GC\_OPTS** to twice that of the heap memory used by PolicySync. You can change the value based on the actual service scenario. Refer to **Step 2** to view the PolicySync heap memory usage.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

**NOTICE**

When the service is rebooted, it becomes unavailable and can disrupt business operations. When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

**Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End



## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

# 7.12.302 ALM-45290 PolicySync Direct Memory Usage Exceeds the Threshold

## Alarm Description

The system checks the direct memory usage of the PolicySync service every 60 seconds. This alarm is generated when the direct memory usage of the PolicySync instance exceeds the threshold (90% of the maximum memory) for five consecutive times. This alarm is cleared when the PolicySync direct memory usage is less than or equal to the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45290	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Direct memory overflow may cause service breakdown.

## Possible Causes

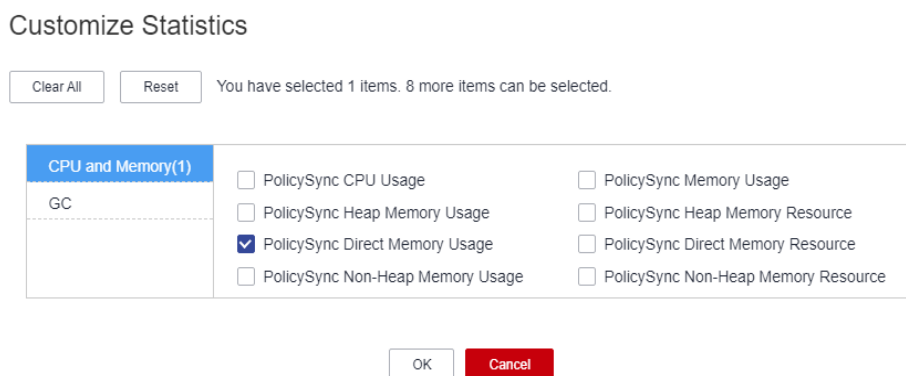
The direct memory of the PolicySync process is overused or the direct memory is inappropriately allocated.

## Handling Procedure

**Check the direct memory usage.**

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms > ALM-45290 PolicySync Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > PolicySync Direct Memory Usage**. Click **OK**.

**Figure 7-177** PolicySync Direct Memory Usage



- Step 3** Check whether the direct memory used by the PolicySync reaches the threshold (90% of the maximum direct memory by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > PolicySync**. Click **Instance Configuration** and then **All Configurations**, and choose **PolicySync > System**. Set **-XX:MaxDirectMemorySize** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

### NOTE

If this alarm is generated, the direct memory configured for PolicySync cannot meet the direct memory required by the PolicySync process. You are advised to check the direct memory usage of PolicySync and change the value of **-XX:MaxDirectMemorySize** in **GC\_OPTS** to the twice of the direct memory used by PolicySync. You can change the value based on the actual service scenario. Refer to **Step 2** to view the TokenServer direct memory usage.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

#### NOTICE

When the service is rebooted, it becomes unavailable and can disrupt business operations. When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

#### Collect fault information.

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

**Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.303 ALM-45291 PolicySync Non-Heap Memory Usage Exceeds the Threshold

### Alarm Description

The system checks the non-heap memory usage of the PolicySync service every 60 seconds. This alarm is generated when the non-heap memory usage of the PolicySync instance exceeds the threshold (90% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45291	Major	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

Non-heap memory overflow may cause service breakdown.

## Possible Causes

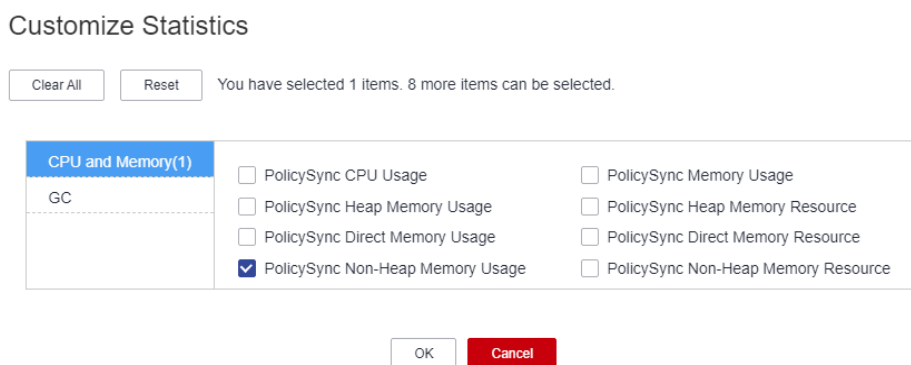
The non-heap memory of the PolicySync instance is overused or the non-heap memory is inappropriately allocated.

## Handling Procedure

**Check non-heap memory usage.**

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms > ALM-45291 PolicySync Non-Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > PolicySync Non-Heap Memory Usage**. Click **OK**.

**Figure 7-178** PolicySync Non-Heap Memory Usage



**Step 3** Check whether the non-heap memory used by PolicySync reaches the threshold (90% of the maximum heap memory by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

**Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > PolicySync**. Click **Instance Configuration** and then **All Configurations**, and choose **PolicySync > System**. Set **-XX:MaxPermSize** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

 **NOTE**

If this alarm is generated, the non-heap memory size configured for the PolicySync instance cannot meet the non-heap memory required by the PolicySync process. You are advised to change the value of **-XX:MaxPermSize** in **GC\_OPTS** to twice that of the current non-heap memory size or change the value based on site requirements.

**Step 5** Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

---

**NOTICE**

When the service is rebooted, it becomes unavailable and can disrupt business operations. When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

---

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

**Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.304 ALM-45292 PolicySync GC Duration Exceeds the Threshold

### Alarm Description

The system checks the GC duration of the PolicySync process every 60 seconds. This alarm is generated when the GC duration of the PolicySync process exceeds the threshold for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45292	Critical (default threshold: 20000ms) Major (default threshold: 12000ms)	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

### Impact on the System

PolicySync responds slowly.

### Possible Causes

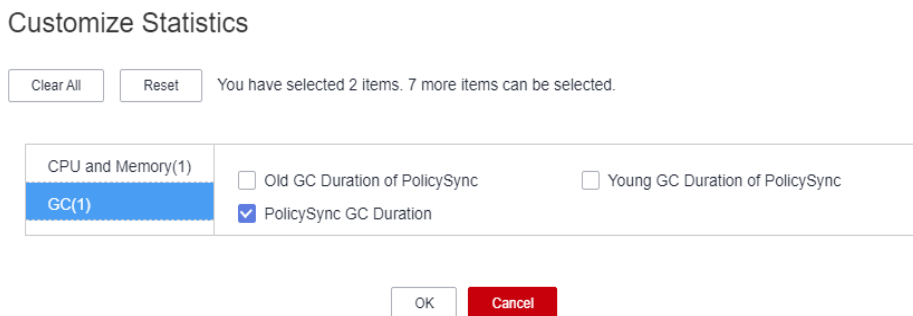
The heap memory of the PolicySync process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

### Handling Procedure

**Check the GC duration.**

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms > ALM-45292 PolicySync GC Duration Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize > GC > PolicySync GC Duration**. Click **OK**.

**Figure 7-179** PolicySync GC Duration



- Step 3** Check whether the GC duration of the PolicySync process collected every minute exceeds the threshold (12 seconds by default).
- If yes, go to **Step 4**.
  - If no, go to **Step 6**.
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Ranger > Instance > PolicySync**. Click **Instance Configuration** and then **All Configurations**, and choose **PolicySync > System**. Set **-Xmx** in the **GC\_OPTS** parameter to a larger value based on site requirements and save the configuration.

**NOTE**

If this alarm is generated, the heap memory configured for PolicySync cannot meet the heap memory required by the PolicySync process. You are advised to change the value of **-Xmx** in **GC\_OPTS** to twice that of the heap memory used by PolicySync. You can change the value based on the actual service scenario. Refer to **Step 2** to view the PolicySync heap memory usage.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 6**.

**NOTICE**

When the service is rebooted, it becomes unavailable and can disrupt business operations. When the instance is rebooted, it cannot be used and any tasks running on the current instance node will fail.

**Collect fault information.**

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 7** Expand the **Service** drop-down list, and select **Ranger** for the target cluster.
- Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None.

## 7.12.305 ALM-45325 Presto Service Unavailable

### NOTE

This section applies only to MRS 3.1.5 or later.

## Alarm Description

The system checks the Presto service status every 60 seconds. This alarm is generated when the Presto service is unavailable. This alarm is cleared when the Presto service recovers.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45325	Critical	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.



## Impact on the System

Presto cannot execute SQL statements.

## Possible Causes

- The Presto coordinator or worker process is faulty.
- The network communication between Presto coordinator and worker instances is interrupted.

## Handling Procedure

**Check the status of the coordinator and worker processes.**

**Step 1** Log in to FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click the **Instance** tab. In the Presto instance list, check whether the running status of all coordinator or worker instances is **Unknown**.

- If yes, go to [2](#).
- If no, go to [4](#).

**Step 2** Above the Presto instance list, click **More** and select **Restart Service** to restart the coordinator and worker processes.


**Step 3** In the alarm list, check whether **ALM-45325 Presto Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [4](#).

**Collect fault information.**

**Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 5** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 6** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 7** Contact the O&M engineers and send the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## 7.12.306 ALM-45326 Number of Presto Coordinator Threads Exceeds the Threshold

### Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45326	Minor	Yes

### Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System

None

### Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

### Handling Procedure

**Check the number of concurrent tasks.**

**Step 1** Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to [2](#).
- If no, go to [Step 4](#).

**Adjust the alarm threshold for the number of threads.**

**Step 2** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or

**Number of Threads (Worker).** Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Upgrade coordinator specifications or add worker node groups.**

**Step 4** Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.307 ALM-45327 Presto Coordinator Process GC Time Exceeds the Threshold

## Alarm Description

The system collects the garbage collection (GC) time of the Presto coordinator process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the coordinator process is less than or equal to the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45327	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the GC time of the coordinator process is too long, the coordinator process performance will be adversely affected, and the coordinator process will even become unavailable.

## Possible Causes

The heap memory of the coordinator process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

### Check the GC duration.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click the **Instance** tab and click the **Coordinator** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.
  - If yes, go to [3](#).
  - If no, go to [6](#).

**Step 3** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Coordinator > JVM**, increase the value of **Xmx** in the **JAVA\_OPTS** parameter, and click **Save**.

**Step 4** Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, select **Presto** for the target cluster, and click **OK**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact the O&M engineers and send the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.308 ALM-45328 Presto Worker Process GC Time Exceeds the Threshold

## Alarm Description

The system collects the garbage collection (GC) time of the Presto worker process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time is less than or equal to the alarm threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45328	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the GC time of the worker process is too long, the worker process performance will be adversely affected, and the worker process will even become unavailable.

## Possible Causes

The heap memory of the worker process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

**Check the GC duration.**


- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.
  - If yes, go to [3](#).
  - If no, go to [6](#).
- Step 3** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker > JVM**, increase the value of **Xmx** in the **JAVA\_OPTS** parameter, and click **Save**.
- Step 4** Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.
- Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to 6.

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

**Step 7** Expand the **Service** drop-down list, select **Presto** for the target cluster, and click **OK**.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.309 ALM-45329 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold

### Alarm Description

The system queries the number of queuing tasks in a resource group through the JMX interface. This alarm is generated when the system detects that the number of queuing tasks in a resource group exceeds the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45329	Minor	Yes

### Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

A large number of tasks may be in the queuing state and cannot be processed as expected. When the number of queuing tasks in a resource group exceeds the threshold (**maxQueued**), new tasks cannot be executed.

## Possible Causes

The resource group configuration is improper or too many tasks in the resource group are submitted.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Coordinator > Customization**, and change the value of **resourceGroupAlarm** in the **resource-groups** parameter to change the threshold for each resource group.

**Step 2** Collect fault information.

1. Log in to the cluster node based on the host name in the fault information and query the number of queuing tasks on the Presto client based on **Resource Group** in the additional information.
2. Log in to the cluster node based on the host name in the fault information, view the **/var/log/Bigdata/nodeagent/monitorlog/monitor.log** file, and search for resource group information to view the monitoring collection information of the resource group.
3. Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None



## 7.12.310 ALM-45330 Number of Presto Worker Threads Exceeds the Threshold

### NOTE

This section applies only to MRS 3.1.5 or later.

### Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45330	Minor	Yes

### Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System

None

### Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

### Handling Procedure

**Check the number of concurrent tasks.**

**Step 1** Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to [2](#).
- If no, go to [4](#).

**Adjust the alarm threshold for the number of threads.**

**Step 2** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Upgrade coordinator specifications or add worker node groups.**

**Step 4** Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.311 ALM-45331 Number of Presto Worker1 Threads Exceeds the Threshold

## Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45331	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

None

## Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

## Handling Procedure

**Check the number of concurrent tasks.**

**Step 1** Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to [2](#).
- If no, go to [4](#).

**Adjust the alarm threshold for the number of threads.**

**Step 2** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Upgrade coordinator specifications or add worker node groups.**

**Step 4** Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

#### Collect fault information.

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.312 ALM-45332 Number of Presto Worker2 Threads Exceeds the Threshold

## Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45332	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

None

## Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

## Handling Procedure

**Check the number of concurrent tasks.**

**Step 1** Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to [2](#).
- If no, go to [4](#).

**Adjust the alarm threshold for the number of threads.**

**Step 2** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Upgrade coordinator specifications or add worker node groups.**

**Step 4** Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.

**Step 5** Check whether the alarm is cleared.


- If yes, no further action is required.

- If no, go to 6.

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.313 ALM-45333 Number of Presto Worker3 Threads Exceeds the Threshold

## Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45333	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

None

## Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

## Handling Procedure

### Check the number of concurrent tasks.

**Step 1** Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to [2](#).
- If no, go to [4](#).

### Adjust the alarm threshold for the number of threads.

**Step 2** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

### Upgrade coordinator specifications or add worker node groups.

**Step 4** Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

### Collect fault information.

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.314 ALM-45334 Number of Presto Worker4 Threads Exceeds the Threshold

## Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45334	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

None

## Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

## Handling Procedure

**Step 1** Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.



- If yes, go to [2](#).
- If no, go to [Step 4](#).

**Adjust the alarm threshold for the number of threads.**

**Step 2** Log in to FusionInsight Manager and choose **O&M > Alarm > Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

**Step 3** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Upgrade coordinator specifications or add worker node groups.**

**Step 4** Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.315 ALM-45335 Presto Worker1 Process GC Time Exceeds the Threshold

## Alarm Description

The system collects the garbage collection (GC) time of the Presto worker1 process every 30 seconds. This alarm is generated when the GC time exceeds the

threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the worker1 process is less than or equal to the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45335	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the GC time of the worker1 process is too long, the worker1 process performance will be adversely affected, and the worker1 process will even become unavailable.

## Possible Causes

The heap memory of the worker1 process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

### Check the GC duration.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker1** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.

- If yes, go to [3](#).
- If no, go to [6](#).

**Step 3** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker1 > JVM**, increase the value of **Xmx** in the **JAVA\_OPTS** parameter, and click **Save**.

**Step 4** Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

#### Collect fault information.

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.316 ALM-45336 Presto Worker2 Process GC Time Exceeds the Threshold

## Alarm Description

The system collects the garbage collection (GC) time of the Presto worker2 process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the worker2 process is less than or equal to the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45336	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the GC time of the worker2 process is too long, the worker2 process performance will be adversely affected, and the worker2 process will even become unavailable.

## Possible Causes

The heap memory of the worker2 process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

**Check the GC duration.**

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker2** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.
  - If yes, go to [3](#).
  - If no, go to [6](#).
- Step 3** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker2 > JVM**, increase the value of **Xmx** in the **JAVA\_OPTS** parameter, and click **Save**.
- Step 4** Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

**Collect fault information.**

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.317 ALM-45337 Presto Worker3 Process GC Time Exceeds the Threshold

## Alarm Description

The system collects the garbage collection (GC) time of the Presto worker3 process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the worker3 process is less than or equal to the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45337	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the GC time of the worker3 process is too long, the worker3 process performance will be adversely affected, and the worker3 process will even become unavailable.

## Possible Causes


The heap memory of the worker3 process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

### Check the GC duration.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.
- Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker3** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.
  - If yes, go to [3](#).
  - If no, go to [6](#).
- Step 3** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker3 > JVM**, increase the value of **Xmx** in the **JAVA\_OPTS** parameter, and click **Save**.
- Step 4** Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.
- Step 5** Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [6](#).

### Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
  - Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.
  - Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
  - Step 9** Contact O&M personnel and provide the collected logs.
- End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.318 ALM-45338 Presto Worker4 Process GC Time Exceeds the Threshold

## Alarm Description

The system collects the garbage collection (GC) time of the Presto worker4 process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the worker4 process is less than or equal to the threshold.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45338	Minor	Yes

## Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

## Impact on the System

If the GC time of the worker4 process is too long, the worker4 process performance will be adversely affected, and the worker4 process will even become unavailable.

## Possible Causes

The heap memory of the worker4 process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

## Handling Procedure

### Check the GC duration.

**Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.

**Step 2** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker4** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.

- If yes, go to [3](#).
- If no, go to [6](#).

**Step 3** On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services > Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker4 > JVM**, increase the value of **Xmx** in the **JAVA\_OPTS** parameter, and click **Save**.

**Step 4** Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.


**Step 5** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [6](#).

### Collect fault information.

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **Presto** for the target cluster.

**Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.



**Step 9** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.319 ALM-45425 ClickHouse Service Unavailable

## Alarm Description

The alarm module checks the ClickHouse instance status every 60 seconds. This alarm is generated when the alarm module detects that all ClickHouse instances are abnormal.

This alarm is cleared when the system detects that any ClickHouse instance is restored and the alarm is cleared.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45425	Critical	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

The ClickHouse service is abnormal. You cannot use FusionInsight Manager to perform cluster operations on the ClickHouse service. The ClickHouse service function is unavailable.

## Possible Causes

The configuration information in the **metrika.xml** file in the component configuration directory of the faulty ClickHouse instance node is inconsistent with that of the corresponding ClickHouse instance in the ZooKeeper.

## Handling Procedure

**Check whether the configuration in metrika.xml of the ClickHouse instance is correct.**

**Step 1** Log in to FusionInsight Manager, choose **Cluster > Services > ClickHouse > Instance**, and locate the abnormal ClickHouse instance based on the alarm information.

- If yes, go to [Step 2](#).
- If no, go to [Step 9](#).

**Step 2** Log in to the host where the ClickHouse service is abnormal and ping the IP address of another normal ClickHouse instance node to check whether the network connection is normal.

- If yes, go to [Step 3](#).
- If no, contact the network administrator to repair the network.

**Step 3** Log in to the node where the client is installed as the client installation user and run the following commands:

```
cd {Client installation path}
```

```
source bigdata_env
```

- For a cluster with Kerberos authentication enabled (security mode):  

```
kinit Component service user
```

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --port 9440 --secure
```
- For a cluster with Kerberos authentication disabled (normal mode):  

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --user Username --password --port 9000
```

Run the following command to query the value of **macros.id**:

```
select substitution from system.macros where macro='id';
```

**Step 4** Log in to the host where the ZooKeeper client is located and log in to the ZooKeeper client.

Switch to the client installation directory.

Example: 

```
cd /opt/client
```

Run the following command to configure environment variables:

```
source bigdata_env
```

Run the following command to authenticate the user (skip this step in common mode):

**kinit** *Component service user*

Run the following command to log in to the client tool:

**zkCli.sh -server** *service IP address of the node where the ZooKeeper role instance locates:client port*

**Step 5** Run the following command to check whether the ClickHouse cluster topology information can be obtained.

**get /clickhouse/config/value of macros.id in Step 3/metrika.xml**

- If yes, go to **Step 6**.
- If no, go to **Step 9**.

**Step 6** Log in to the host where the ClickHouse instance is abnormal and go to the configuration directory of the ClickHouse instance.

```
cd ${BIGDATA_HOME}/FusionInsight_ClickHouse_Version/
x_x_ClickHouseServer/etc
```

```
cat metrika.xml
```

**Step 7** Check whether the cluster topology information on ZooKeeper obtained in **Step 5** is the same as that in the **metrika.xml** file in the component configuration directory in **Step 6**.

- If yes, check whether the alarm is cleared. If the alarm persists, go to **Step 9**.
- If no, go to **Step 8**.

**Step 8** On FusionInsight Manager, choose **Cluster > Services > ClickHouse**, click **More**, and select **Synchronize Configuration**. Then, check whether the service status is normal and whether the alarm is cleared 5 minutes later.


- If yes, no further action is required.
- If no, go to **Step 9**.

**Collect the fault information.**

**Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 10** Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

**Step 11** Choose the corresponding host form the host list.

**Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 13** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

# 7.12.320 ALM-45426 ClickHouse Service Quantity Quota Usage in ZooKeeper Exceeds the Threshold

## Alarm Description

The alarm module checks the quota usage of the ClickHouse service in the ZooKeeper every 60 seconds. This alarm is generated when the alarm module detects that the usage exceeds the threshold (90%).

This alarm is cleared when the system detects that the usage is lower than the threshold and the alarm is cleared.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45426	Major (default)	Yes

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

After the ZooKeeper quantity quota of the ClickHouse service exceeds the threshold, you cannot perform cluster operations on the ClickHouse service on FusionInsight Manager. As a result, the ClickHouse service cannot be used.

## Possible Causes

- When table data is created, inserted, or deleted, the ClickHouse creates znodes on ZooKeeper nodes. As the service volume increases, the number of znodes may exceed the configured threshold.

- No quota limit is set for the metadata directory **/clickhouse** of ClickHouse in ZooKeeper.

## Handling Procedure

### Check the number of znodes created by ClickHouse on ZooKeeper.

**Step 1** Log in to the host where the ZooKeeper client is located and log in to the ZooKeeper client.

Switch to the client installation directory.

Example: `cd /opt/client`

Run the following command to configure environment variables:

`source bigdata_env`

Run the following command to authenticate the user (skip this step in common mode):

`kinit Component service user`

Run the following command to log in to the client tool:

`zkCli.sh -server service IP address of the node where the ZooKeeper role instance locates:client port`

**Step 2** Run the following command to check the quota used by the ClickHouse in the ZooKeeper and check whether the quota information is correctly set:

`listquota /clickhouse`

```
absolute path is /zookeeper/quota/clickhouse
Quota for path /clickhouse does not exist.
```

If the preceding information indicates that the quota configuration is incorrect, go to [Step 3](#).

If no, go to [Step 5](#).

**Step 3** Log in to FusionInsight Manager and choose **Cluster > Services > ZooKeeper**. On the displayed page, click **Configurations** and click **All Configurations**. On this sub-tab page, search for **quotas.auto.check.enable** to check whether its value is **true**.

If the value is not **true**, change the value to **true** and click **Save**.

**Step 4** On FusionInsight Manager, choose **Cluster > Services > ClickHouse**, click **More**, and select **Synchronize Configuration**. After the synchronization is successful, go to [Step 1](#).

**Step 5** Run the following command and check whether the ratio of the **count** value of **Output stat** to the **count** value of **Output quota** in the command output is greater than **0.9**:

`listquota /clickhouse`

```
absolute path is /zookeeper/quota/clickhouse
Output quota for /clickhouse count=200000,bytes=1000000000
Output stat for /clickhouse count=2667,bytes=60063
```

In the preceding information, the **count** value of **Output stat** is **2667**, and the **count** value of **Output quota** is **200000**.

- If yes, go to [Step 6](#).
- If no, check whether the alarm is cleared 5 minutes later. If the alarm persists, go to [Step 8](#).

**Step 6** On FusionInsight Manager, choose **Cluster > Services > ClickHouse**. Click **Configurations** and then **All Configurations**. Search for **clickhouse.zookeeper.quota.node.count**, set it to a value twice the **count** of **Output stat** in [Step 5](#). **Do not use a value larger than 6000000. Otherwise, there will be high risks. Exercise caution when setting this parameter.**

**Step 7** Restart the ClickHouse instance for which the alarm is generated, and check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, perform [Step 6](#) again, and check whether the alarm is cleared 5 minutes later. If the alarm persists, go to [Step 8](#).

---

**NOTICE**

During the instance restart, the instance is unavailable, and the ClickHouse service on that instance fails to be executed.


---

**Collect the fault information.**

**Step 8** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 9** Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

**Step 10** Choose the corresponding host from the host list.

**Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.321 ALM-45427 ClickHouse Service Capacity Quota Usage in ZooKeeper Exceeds the Threshold

### Alarm Description

The alarm module checks the quota usage of the ClickHouse service in the ZooKeeper every 60 seconds. This alarm is generated when the alarm module detects that the usage exceeds the threshold (90%).

This alarm is cleared when the system detects that the usage is lower than the threshold and the alarm is cleared.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45427	Major (default)	Yes

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System

After the ZooKeeper quantity quota of the ClickHouse service exceeds the threshold, you cannot perform cluster operations on the ClickHouse service on FusionInsight Manager. As a result, the ClickHouse service cannot be used.

### Possible Causes

- When table data is created, inserted, or deleted, the ClickHouse creates znodes on ZooKeeper nodes. As the service volume increases, the capacity of znodes may exceed the configured threshold.
- No quota limit is set for the metadata directory **/clickhouse** of ClickHouse in ZooKeeper.

## Handling Procedure

### Check the znode capacity of the ClickHouse in the ZooKeeper.

**Step 1** Log in to the host where the ZooKeeper client is located and log in to the ZooKeeper client.

Switch to the client installation directory.

Example: `cd /opt/client`

Run the following command to configure environment variables:

`source bigdata_env`

Run the following command to authenticate the user (skip this step in common mode):

`kinit Component service user`

Run the following command to log in to the client tool:

`zkCli.sh -server service IP address of the node where the ZooKeeper role instance locates:client port`

**Step 2** Run the following command to check the quota used by the ClickHouse in the ZooKeeper and check whether the quota information is correctly set:

`listquota /clickhouse`

```
absolute path is /zookeeper/quota/clickhouse
Quota for path /clickhouse does not exist.
```

- If the preceding information indicates that the quota configuration is incorrect, go to [Step 3](#).
- If not, go to [Step 5](#).

**Step 3** Log in to FusionInsight Manager and choose **Cluster > Services > ZooKeeper**. On the displayed page, click **Configurations** and click **All Configurations**. On this sub-tab page, search for **quotas.auto.check.enable** to check whether its value is **true**.

If the value is not **true**, change the value to **true** and click **Save**.

**Step 4** On FusionInsight Manager, choose **Cluster > Services > ClickHouse**, click **More**, and select **Synchronize Configuration**. After the synchronization is successful, go to [Step 1](#).

**Step 5** Run the following command and check whether the ratio of the **bytes** value of **Output stat** to the **bytes** value of **Output quota** in the command output is greater than **0.9**:

`listquota /clickhouse`

```
absolute path is /zookeeper/quota/clickhouse
Output quota for /clickhouse count=200000,bytes=1000000000
Output stat for /clickhouse count=2667,bytes=60063
```

In the preceding information, the **bytes** value of **Output stat** is **60063**, and the **bytes** value of **Output quota** is **1000000000**.

- If yes, go to [Step 6](#).
- If no, check whether the alarm is cleared 5 minutes later. If the alarm persists, go to [Step 8](#).



- Step 6** On FusionInsight Manager, choose **Cluster > Services > ClickHouse**. Click **Configurations** and then **All Configurations**. Search for **clickhouse.zookeeper.quota.size**, set it to a value twice the **bytes** of **Output stat** in **Step 5**. **Do not use a value larger than 6000000. Otherwise, there will be high risks. Exercise caution when setting this parameter.**
- Step 7** Restart the ClickHouse instance for which the alarm is generated, and check whether the alarm is cleared 5 minutes later.
- If yes, no further action is required.
  - If no, perform **Step 6** again, and check whether the alarm is cleared 5 minutes later. If the alarm persists, go to **Step 8**.


---

**NOTICE**

During the instance restart, the instance is unavailable, and the ClickHouse service on that instance fails to be executed.

---

**Collect fault information.**

- Step 8** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 9** Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.
- Step 10** Choose the corresponding host from the host list.
- Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 12** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.322 ALM-45428 ClickHouse Disk I/O Exception

### Alarm Description

This alarm is generated when the alarm module detects EIO or EROFS errors during ClickHouse read and write every 60 seconds.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45428	Major (default)	No

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

- ClickHouse fails to read and write data. The INSERT, SELECT, and CREATE operations on the local tables may be abnormal. Distributed tables are not affected.
- Services are affected, and I/Os fail.

## Possible Causes

The disk is aged or has bad sectors.

## Handling Procedure

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45428 ClickHouse Disk I/O Exception**. Check the role name and the IP address of the host where the alarm is generated in **Location**.
- Step 2** Use PuTTY to log in to the node for which the fault is generated as user **root**.
- Step 3** Run the **df -h** command to check the mount directory and find the disk mounted to the faulty directory.
- Step 4** Run the **smartctl -a /dev/sdFaulty disk** command to check the disk. In the command, *Faulty disk* indicates the disk obtained in [Step 3](#).
  - If **SMART Health Status: OK** is displayed, as shown in the following figure, the disk is healthy. In this case, go to [Step 6](#).

```
=== START OF READ SMART DATA SECTION ===
SMART Health Status: OK

Current Drive Temperature: 26 C
Drive Trip Temperature: 60 C

Manufactured in week 50 of year 2018
Specified cycle count over device lifetime: 10000
Accumulated start-stop cycles: 25
Specified load-unload count over device lifetime: 300000
Accumulated load-unload cycles: 356
Elements in grown defect list: 0
```

- If the number following **Elements in grown defect list** is not 0, as shown in the following figure, the disk may have bad sectors. If **SMART Health Status: FAILURE** is displayed, the disk is in the sub-health state. In this case, contact O&M personnel.

```
=== START OF READ SMART DATA SECTION ===
SMART Health Status: FAILURE PREDICTION THRESHOLD EXCEEDED: ascq=0x5 [asc=5d, ascq=5]

Current Drive Temperature: 30 C
Drive Trip Temperature: 60 C

Manufactured in week 50 of year 2018
Specified cycle count over device lifetime: 10000
Accumulated start-stop cycles: 28
Specified load-unload count over device lifetime: 300000
Accumulated load-unload cycles: 354
Elements in grown defect list: 5344
```

**Step 5** After the fault is rectified, manually clear the alarm on FusionInsight Manager and check whether the alarm is generated again during the periodic check.


- If yes, go to [Step 6](#).
- If no, no further action is required.

#### Collect fault information.

**Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 7** Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

**Step 8** Choose the corresponding host from the host list.

**Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

**Step 10** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

If the alarm has no impact, manually clear the alarm.

## Related Information

None

## 7.12.323 ALM-45429 Table Metadata Synchronization Failed on the Added ClickHouse Node

### NOTE

This section applies only to MRS 3.1.2 or later.

### Alarm Description

This alarm is generated when the local table corresponding to the distributed table fails to be created during ClickHouse capacity expansion.

### Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45429	Major	No

### Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

### Impact on the System

The distributed table fails to be queried.

### Possible Causes

A node is stopped or faulty during capacity expansion.

### Handling Procedure

**Step 1** On FusionInsight Manager, choose **Cluster > Services > ClickHouse > Instance**.

**Step 2** Check whether an instance is stopped, decommissioned, or faulty.

- If yes, go to [Step 3](#).

- If no, go to [Step 4](#).

**Step 3** Start the instance or rectify the instance fault until all instances are running properly.

**Step 4** On FusionInsight Manager, choose **O&M > Alarm > Alarms**, locate this alarm and the faulty host based on the location information.

**Step 5** Log in to the faulty host as user **omm**.

**Step 6** Run the following commands to initialize environment variables:

```
source Cluster installation directory/FusionInsight_ClickHouse_*/
*_*_ClickHouseServer/etc/ENV_VARS
```

```
source Cluster installation directory/FusionInsight_ClickHouse_*/
*_*_ClickHouseServer/etc/clickhouse-env.sh
```

```
export CLICKHOUSE_CONF_DIR=${CLICKHOUSE_CONF_DIR}
```

**Step 7** Run the following command to run the metadata synchronization tool to synchronize metadata from the existing node to the faulty node:

```
sh Cluster installation directory/FusionInsight_ClickHouse_*/install/
FusionInsight-ClickHouse-*/clickhouse/sbin/clickhouse-create-meta.sh true
```

**Step 8** Run the following command to view the log information and check whether the metadata has been synchronized:

```
vim /var/log/Bigdata/clickhouse/clickhouseServer/start.log
```

- If the synchronization is complete, go to [Step 9](#).
- If the synchronization fails, go to [Step 10](#).


**Step 9** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the **Alarm ID** column, locate the corresponding alarm and click **Clear** in the **Operation** column. In the displayed dialog box, click **OK** to manually clear the alarm.

**Collect fault information.**

**Step 10** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

**Step 11** Expand the **Service** drop-down list, select **ClickHouse** for the target cluster, and click **OK**.

**Step 12** Choose the corresponding host from the host list.

**Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

**Step 14** Contact O&M personnel and provide the collected logs.

----End

## Alarm Clearance

This alarm needs to be manually cleared after the fault is rectified.

## Related Information

None

### 7.12.324 ALM-45430 Permission Metadata Synchronization Failed on the Added ClickHouse Node

#### NOTE

This section applies only to MRS 3.1.2 or later.

## Alarm Description

This alarm is generated when user and permission information fails to be synchronized during ClickHouse capacity expansion.

## Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45430	Major	No

## Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

## Impact on the System

The created user does not have operation permissions on the node.

## Possible Causes


A node is stopped or faulty during capacity expansion.

## Handling Procedure

**Step 1** On FusionInsight Manager, choose **Cluster > Services > ClickHouse > Instance**.

- Step 2** Check whether an instance is stopped, decommissioned, or faulty.
- If yes, go to [Step 3](#).
  - If no, go to [Step 4](#).
- Step 3** Start the instance or rectify the instance fault until all instances are running properly.
- Step 4** On FusionInsight Manager, choose **O&M > Alarm > Alarms**, locate this alarm and the faulty host based on the location information.
- Step 5** Log in to the faulty host as user **omm**.
- Step 6** Run the following commands to initialize environment variables:
- ```
source ${BIGDATA_HOME}/FusionInsight_ClickHouse_*/
*_ClickHouseServer/etc/ENV_VARS

source ${BIGDATA_HOME}/FusionInsight_ClickHouse_*/
*_ClickHouseServer/etc/clickhouse-env.sh

export CLICKHOUSE_CONF_DIR=${CLICKHOUSE_CONF_DIR}
```
- Step 7** Run the following command to run the metadata synchronization tool to synchronize metadata from the existing node to the faulty node:
- ```
sh ${BIGDATA_HOME}/FusionInsight_ClickHouse_*/install/FusionInsight-
ClickHouse-*/clickhouse/sbin/clickhouse-create-meta.sh true
```
- Step 8** Run the following command to view the log information and check whether the metadata has been synchronized:
- ```
vim /var/log/Bigdata/clickhouse/clickhouseServer/start.log
```
- If the synchronization is complete, go to [Step 9](#).
- If the synchronization fails, go to [Step 10](#).
- Step 9** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the **Alarm ID** column, locate the corresponding alarm and click **Clear** in the **Operation** column. In the displayed dialog box, click **OK** to manually clear the alarm.
- Collect the fault information.**
- Step 10** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 11** Expand the **Service** drop-down list, select **ClickHouse** for the target cluster, and click **OK**.
- Step 12** Choose the corresponding host from the host list.
- Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 14** Contact O&M personnel and provide the collected logs.
- End

Alarm Clearance

This alarm needs to be manually cleared after the fault is rectified.

Related Information

None

7.12.325 ALM-45431 Improper ClickHouse Instance Distribution for Topology Allocation

Alarm Description

The ClickHouseServer instance distribution does not meet the topology allocation requirements.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 45431 | Critical | No |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |

Impact on the System

Some ClickHouseServer instances are unavailable.

Possible Causes

During installation or capacity expansion, the number of instances or allocation mode does not meet the topology requirements.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**, locate the row that contains the alarm, and analyze the cause based on **Location** and **Additional Information**.

Step 2 Handle the alarm based on the additional alarm information and handling method in the following table.

| Additional Information | Remarks | Handling Method |
|---|--|--|
| <p><i>n</i> ClickHouseServer instances should be added to other AZ.</p> | <p>This alarm is generated when a single cluster is deployed in cross-AZ DR mode. The ClickHouseServer instance deployment does not meet the cross-AZ DR topology allocation requirements. As a result, some instances cannot work properly.</p> | <ol style="list-style-type: none"> 1. On FusionInsight Manager, choose Cluster > Services > ClickHouse, click the Instance tab, locate the row that contains the alarm, view the host name in Location, and find the AZ for which the alarm is generated in the AZ column based on the host name. 2. On the Instance page, click Add Instance to add <i>n</i> ClickHouseServer instances to other AZs except the AZ where the alarm is generated. |
| <p><i>n</i> ClickHouseServer instances should be added.</p> | <p>This alarm is generated when a non-single-cluster is deployed in the default deployment mode of cross-AZ DR. The number of ClickHouseServer instances in the cluster is less than an even number. As a result, some instances cannot work properly.</p> | <ol style="list-style-type: none"> 1. Determine the number (<i>n</i>) of ClickHouseServer instances to be added based on the alarm information. 2. On FusionInsight Manager, choose Cluster > Services > ClickHouse, click the Instance tab, and add <i>n</i> ClickHouseServer instances to the cluster. |

----End

Alarm Clearance

This alarm needs to be manually cleared after the fault is rectified.

Related Information

None

7.12.326 ALM-45432 ClickHouse User Synchronization Process Fails

Alarm Description

The system checks the status of the ClickHouse user role synchronization process every 5 minutes. This alarm is generated when the system detects that the ClickHouse user role synchronization process is faulty or the user role synchronization fails.

This alarm is automatically cleared when the ClickHouse user role synchronization process or function becomes normal.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 45432 | Major | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| RoleName | Specifies the role for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |

Impact on the System

Some ClickHouseServer instances are unavailable.

Possible Causes

- The ClickHouse user role synchronization process is not started properly or exits abnormally.

- The ClickHouse user role synchronization process fails to synchronize user role information because the LdapServer service is faulty.

Handling Procedure

Check whether the ClickHouse user role synchronization process is normal.

Step 1 Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. On the page that is displayed, search for **ALM-45432 ClickHouse User Synchronization Process Fails**.

Step 2 Check the host name and additional information in the alarm details.

- If the additional information is "Process clickhouse-ugsync is not exit.", go to [Step 3](#).
- If the additional information is "Process clickhouse-ugsync sync user failed.", go to [Step 6](#).

Step 3 Log in to the faulty host as user **omm** and run the following command to check whether the ClickHouse user role synchronization process is normal:

```
ps -ef | grep 'clickhouse-ugsync'
```

Abnormal result of the synchronization process:

```
[omm@server-2110081635-0001 ~]$ ps -ef | grep 'clickhouse-ugsync'
omm      20104 13146 0 15:57 pts/7    00:00:00 grep --color=auto clickhouse-ugsync
```

- If yes, the alarm is automatically cleared. If the alarm is cleared, no further action is required. If the alarm persists, go to [Step 8](#).
- If no, go to [Step 4](#).

Step 4 Log in to the faulty host as user **omm** and run the following command to check whether the crontab daemon task is correctly configured:

```
crontab -l
```

Normal setting of the crontab daemon task:

```
*/5 * * * * bash /xxxxx/clickhouse_ugsync_check.sh >/dev/null 2>&1
```

- If yes, check whether the alarm is cleared 5 minutes later. If the alarm is cleared, no further action is required. If the alarm persists, go to [Step 8](#).
- If no, go to [Step 5](#).

Step 5 Log in to FusionInsight Manager, choose **Cluster > Services > ClickHouse**. On the page that is displayed, click the **Instance** tab. On this tab page, find the abnormal ClickHouseServer instance based on the fault information, and restart it. Wait for 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

NOTICE

During the instance restart, the instance is unavailable, and the ClickHouse service on that instance fails to be executed.

Check whether the LdapServer service is normal.

Step 6 Log in to FusionInsight Manager, choose **Cluster > Services**, and check whether **Running Status** of LdapServer is **Normal**.

- If yes, go to [Step 8](#).
- If no, go to [Step 7](#).

Step 7 Handle the LdapServer service unavailable alarm according to ALM-25000 LdapServer Service Unavailable.

After **Running Status** of LdapServer becomes **Normal**, check whether this alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 8](#).

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M > Log > Download**.

Step 9 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **ClickHouseServer** for the target cluster.

Step 10 Expand the **Hosts** list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.327 ALM-45433 ClickHouse AZ Topology Exception

Alarm Description

If the cross-AZ HA function is enabled for a cluster where ClickHouse has been deployed, the ClickHouse topology remains unchanged. This alarm is generated when the cross-AZ HA does not take effect if backup nodes of the same shard are in the same AZ.

This alarm is automatically cleared when the system detects that all shards meet the cross-AZ HA deployment requirements.

Alarm Attributes

| Alarm ID | Alarm Severity | Auto Cleared |
|----------|----------------|--------------|
| 45433 | Critical | Yes |

Alarm Parameters

| Parameter | Description |
|-------------|--|
| Source | Specifies the cluster for which the alarm was generated. |
| ServiceName | Specifies the service for which the alarm was generated. |
| HostName | Specifies the host for which the alarm was generated. |

Impact on the System

The current deployment of the ClickHouse service does not support cross-AZ HA.

Possible Causes

After cross-AZ HA is enabled, all backup nodes of a shard are in the same AZ.

Handling Procedure

Modify the AZ of backup nodes.

Step 1 Log in to the node where the client is installed as the client installation user. Run the following command to switch to the client installation directory:

```
cd {Client installation path}
```

Step 2 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 3 Run the following command to authenticate the user (skip this step in normal mode):

```
kinit Component service user
```

Step 4 Run the following command to log in to the client tool:

```
zkCli.sh -server Service IP address of the node where the ZooKeeper instance resides:Client port
```

Step 5 Run the following command to view the current topology:

```
get /clickhouse/topo
```


 **NOTE**

If the ClickHouse is installed with multiple services, run the **get /clickhouse{-n}/topo** command. For example, if the ClickHouse-1 is installed, run the **get /clickhouse-1/topo** command.

```
[zk: 192.168.20.36:24002(CONNECTED) 0] get /clickhouse/topo
<topo>
<mcluster>
<shard id="14" index="1">
<server id="15">
<replica>1</replica>
<az>AZ1</az>
<host>192-168-20-205</host>
<port>21427</port>
</server>
<server id="16">
<replica>2</replica>
<az>AZ1</az>
<host>192-168-20-2205</host>
<port>21427</port>
</server>
</shard>
</mcluster>
</topo>
```

- Step 6** Select a host from the desired shard and deploy the host in another AZ.
- Step 7** Log in to FusionInsight Manager, click **Host**, select the host you have deployed in **Step 6** and choose **More > Reinstall** to reinstall the host.
- Step 8** Choose **Cluster > Cross-AZ HA**, click **Configure AZ and Policy** and change the AZ information of the reinstalled host to the AZ planned in the **Step 6**.
- Step 9** Wait for five minutes and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to **Step 10**.

Collect fault information.

- Step 10** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 11** Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **ClickHouseServer** for the target cluster.
- Step 12** Expand the **Hosts** list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.
- Step 13** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 14** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.328 ALM-45434 A Single Replica Exists in the ClickHouse Data Table

Description

This alarm is generated when a single replica is detected in a custom logical cluster after the custom logical cluster is enabled for ClickHouse.

This alarm is automatically cleared when the system detects that the custom logical cluster uses multiple replicas.

Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 45434 | Major | Yes |

Parameters

| Name | Meaning |
|-------------|---|
| Source | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName | Specifies the role for which the alarm is generated. |
| HostName | Specifies the host for which the alarm is generated. |

Impact on the System


If a hardware fault occurs, data cannot be restored.

Possible Causes

The **metrika.xml** file in the ClickHouse configuration directory contains single-replica configuration.

Procedure

Check whether the configuration in metrika.xml of the ClickHouse instance is correct.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click  to view the name of the host for which the alarm is generated. On the **Hosts** page, view the host IP address based on the host name.

Step 2 Log in to the host where the ClickHouse instance is abnormal, go to the configuration directory of the ClickHouse instance, and run the following commands:

```
cd ${BIGDATA_HOME}/FusionInsight_ClickHouse_Version/  
x_x_ClickHouseServer/etc
```

```
cat metrika.xml
```

Step 3 View the number of shards in each custom logical cluster and check that a single replica exists. Then, go to [Step 4](#).

NOTE

If a shard contains only one node, a single replica exists in a logical cluster, as shown in the following:


```
<shard>  
  <internal_replication>true</internal_replication>  
  <replica>  
    <host>host-name1</host>  
    <port>port</port>  
    <user>clickhouse</user>  
    <password/>  
  </replica>  
</shard>
```

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 5 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 6 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.329 ALM-45435 Inconsistent Metadata of ClickHouse Tables

Alarm Description

This alarm is generated when the metadata in a distributed table or in the local table of the distributed table has been inconsistent for 180 min.

This alarm is automatically cleared when the metadata in the distributed table or in the local table of the distributed table becomes consistent.

Metadata consistency includes:

- Consistent quantity, name, sequence, and type of each column in the table
- Consistent partition keys
- Consistent sorting keys
- Consistent primary keys
- Consistent sampling keys

NOTE

If this alarm exists, table metadata is inconsistent in the ClickHouse cluster to which the current node belongs. The inconsistency may be caused by multiple reasons, not limited to those mentioned in additional information.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45435	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Table	Specifies the database name and table name for which the alarm is generated.

Impact on the System

Subsequent operations such as INSERT and ALTER on the table may fail.

Possible Causes

Table metadata modification fails or is not executed on one or more ClickHouseServer nodes.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

```
cd {Client installation path}
```

```
source bigdata_env
```

- For a cluster with Kerberos authentication enabled (security mode):

```
kinit Component service user
```

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --port 9440 --secure
```

- For a cluster with Kerberos authentication disabled (normal mode):

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --user Username --password --port 9000
```

Step 3 Check whether any task is being executed for the table to which the alarm is generated.

Run the following command to check whether any SQL task is being executed:

```
select * from system.processes where current_database='Database name' and query like '%Table name%'
```

Run the following command to check whether a mutation task is being executed:

```
select * from system.mutations where database='Database name' and table='Table name';
```

- If the query result is empty, go to [Step 4](#).
- If the query result contains error information, rectify the fault accordingly. If the fault cannot be rectified based on the error information, go to [Step 6](#).
- If the query result contains information about an on-going task with no error, the SQL/mutation task is being executed.

Wait for 5 minutes. If the alarm is cleared, no further action is required. If the alarm persists, go to [Step 4](#).

Step 4 Modify the table structure, delete a table, or add a table based on service requirements until the table metadata of all nodes in the cluster is consistent. After 5 minutes, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Step 5 If the table has been deleted, manually clear the alarm and check whether the alarm is reported again.

- If yes, go to [Step 6](#).
- If no, no further action is required.

Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.
- Step 7** Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.
- Step 8** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.
- Step 9** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.330 ALM-45436 Skew ClickHouse Table Data

Alarm Description

This alarm is generated when data skew occurs in the local table of a distributed table between ClickHouse nodes. This alarm is automatically cleared when data becomes balanced.

Data skew check method:

- If **min_table_check_data_bytes** is set to **0**, data skew check is disabled.
- If **min_table_check_data_bytes** is greater than **0**, data skew check is enabled.

After data skew check is enabled, if the data volume in a table is less than the **min_table_check_data_bytes** value, no alarm will be reported due to data skew. When the data volume is greater than the **min_table_check_data_bytes** value and the data volume difference between the same table on different nodes is greater than the percentage specified in **min_table_data_varies_rate**, data skew occurs and this alarm is reported.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45436	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Table	Specifies the database name and table name for which the alarm is generated.

Impact on the System

SQL execution efficiency may be lowered.

Possible Causes

The data write policy is improper, causing unbalanced data among nodes.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

```
cd {Client installation path}
```

```
source bigdata_env
```

- Security mode (with Kerberos enabled):

```
kinit Component service user
```

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --port 9440 --secure
```

- Normal mode (with Kerberos disabled):

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --user Username --password --port 9000
```

Step 3 View data distribution.

```
select FQDN(), database, table, sum(data_compressed_bytes) from clusterAllReplicas(Name of the logical cluster, system.parts) where database='Database name' and table='Table name' and active=1 group by (FQDN(), database, table);
```

Step 4 Balance data with a few clicks or migrate data based on service requirements.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 7 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 8 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.331 ALM-45437 Excessive Parts in the ClickHouse Table

Alarm Description

This alarm is generated when the number of parts exceeds the threshold specified by **part_num_threshold**.

This alarm is automatically cleared when the number of parts is less than the **part_num_threshold** value.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45437	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Table	Specifies the database name and table name for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Service errors may occur.

Possible Causes

The data distribution in the ClickHouse table is improper, or the background merge task is executed slowly.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

```
cd {Client installation path}
```

```
source bigdata_env
```

- Security mode (with Kerberos enabled):

```
kinit Component service user
```

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --port 9440 --secure
```

- Normal mode (with Kerberos disabled):

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --user Username --password --port 9000
```

Step 3 Run the following command to manually merge parts:

```
optimize table Database name.Table name final;
```

Step 4 Check whether the number of parts has decreased.

```
select FQDN(), database, table, count(1) from clusterAllReplicas(default_cluster, system.parts) where database='Database name' and table='Table name' and active=1 group by (FQDN(), database, table);
```

1. If the number of parts is less than the threshold, wait for 5 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 5](#).
2. If the number of parts does not decrease, check whether the partition key of the table is set properly. If the number of partitions is too large, rectify the service logic.
3. If the command output is empty, the table does not exist. This alarm is a historical alarm and can be ignored. Manually clear it.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 7 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.332 ALM-45438 ClickHouse Disk Usage Exceeds 80%

Alarm Description

The system checks the disk capacity of the ClickHouseServer node every 1 minute. This alarm is generated when the usage of the disk where the ClickHouse data directory or metadata directory resides exceeds 80%.

This alarm is automatically cleared when the usage of the disk where the ClickHouse data directory or metadata directory is located is lower than 80%.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45438	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DiskPath	Specifies the path of the disk for which the alarm is generated.

Impact on the System

The ClickHouse write operation may fail.

Possible Causes

The disk capacity of the ClickHouseServer node is too small.

Handling Procedure

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.
- Step 2** Expand the disk capacity of the node for which the alarm is generated.
- Step 3** Go to **Step 4** if the expansion fails or the alarm persists after the expansion.
Collect fault information.
- Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 5** Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.
- Step 6** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.
- Step 7** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 8** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.333 ALM-45439 ClickHouse Node Enters the Read-Only Mode

Alarm Description

The system checks the disk capacity of the ClickHouseServer node every 1 minute. This alarm is generated when the system detects that the disk capacity exceeds 90% and the ClickHouseServer node enters the read-only mode.

This alarm is automatically cleared when the system detects that the disk capacity is lower than 90% and the ClickHouseServer node exits the read-only mode.

NOTE

If the ClickHouseServer node is in read-only mode and you need to log in to the client to clear data, you can manually exit the read-only mode using the following method:

Log in to FusionInsight Manager, choose **Cluster > Services > ClickHouse > Configurations > All Configurations**, search for **profiles.default.readonly**, and change its value to **0**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45439	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DiskPath	Specifies the path of the disk for which the alarm is generated.

Impact on the System

After the ClickHouseServer node enters the read-only mode, all write, modification, and deletion operations fail.

Possible Causes

The disk usage of the ClickHouse node exceeds 90%.

Handling Procedure

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.
- Step 2** Expand the disk capacity of the node for which the alarm is generated.
- Step 3** Go to [Step 4](#) if the expansion fails or the alarm persists after the expansion.

NOTE

After the capacity expansion, this alarm can be automatically cleared only when **profiles.default.readonly** is **auto**. If its value has been manually changed, change it back to **auto**. If **profiles.default.readonly** needs to be set to **0** or **1** based on service requirements, manually clear this alarm.

Collect fault information.

- Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 5** Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.
- Step 6** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.
- Step 7** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 8** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.334 ALM-45440 Inconsistency Between ClickHouse Replicas

Alarm Description

When the number of ClickHouse replicas is greater than 1, the system periodically checks the replicated table. This alarm is generated if replicated table data is not

synchronized. This alarm is cleared when data in all replicated tables between replicas becomes synchronized.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45440	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Table	Specifies the table name for which the alarm is generated.

Impact on the System

The data reliability of the ClickHouse replicated table is affected, causing data differences and affecting the query result of the distributed table.

Possible Causes

- The ClickHouse service is overloaded.
- The connection between the ClickHouse and ZooKeeper is abnormal.

Handling Procedure

Check whether the ClickHouse service load is heavy.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the database name, table name, role name and IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

```
cd {Client installation path}
```

```
source bigdata_env
```

- For a cluster with Kerberos authentication enabled (security mode):
kinit *Component service user*
clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**
- For a cluster with Kerberos authentication disabled (normal mode):
clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--user** *Username* **--password** **--port** 9000

Step 3 Run the following statement to check whether data is frequently written to the system table. If yes, wait until the service execution is complete and check whether the alarm is cleared.

```
SELECT query_id, user, FQDN(), elapsed, query FROM system.processes ORDER BY query_id;
```

- If yes, no further action is required.
- If no, go to [Step 4](#).

Step 4 Check whether a large amount of data is written. If yes, wait until the task is complete and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Step 5 Run the following statement to check whether replicas are synchronized:

```
select table,absolute_delay, queue_size, inserts_in_queue, merges_in_queue from system.replicas where absolute_delay > 0 order by absolute_delay desc limit 10;
```

- If yes, go to [Step 6](#).
- If no, go to [Step 9](#).

Step 6 If **inserts_in_queue** contains a large amount of content to be inserted, run the following SQL statement to query the replica synchronization queue and locate the error cause:

```
SELECT database,table,type,any(last_exception),any(postpone_reason),min(create_time),max(last_attempt_time),max(last_postpone_time),max(num_postponed) AS max_postponed,max(num_tries) AS max_tries,min(num_tries) AS min_tries,countIf(last_exception != "") AS count_err,countIf(num_postponed > 0) AS count_postponed,countIf(is_currently_executing) AS count_executing,count() AS count_all FROM system.replication_queue GROUP BY database,table,type ORDER BY count_all DESC
```

Check whether an error message similar to the following is displayed:

```
Not executing fetch of part xxx because n fetches already executing, max n
```

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

Step 7 On FusionInsight Manager, choose **Cluster > Services > ClickHouse > Configurations > All Configurations**, and check whether the value of **background_pool_size** is twice the number of cores on the node.

- If yes, go to [Step 9](#).
- If no, go to [Step 8](#).

Step 8 Set this parameter to twice the number of cores on the node and synchronize the configuration. Wait for a while and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Check the connectivity between ClickHouse and ZooKeeper.

Step 9 Log in to the node where the ClickHouseServer instance is located, go to `$(BIGDATA_HOME)/FusionInsight_ClickHouse_*/*_ClickHouseServer/etc`, and check whether the port numbers of the ClickHouseServer and ZooKeeper in the `config.xml` file are the same, as shown in the following information in bold:

NOTE

To view the ZooKeeper port number, choose **Cluster > Services > ZooKeeper > Configurations > All Configurations** on FusionInsight Manager, and check the value of `clientPort`.

```
<zookeeper>
<session_timeout_ms>10000</session_timeout_ms>
<node index="1">
  <host>server-2110082001-0019</host>
  <port>24002</port>
</node>
<node index="2">
  <host>server-2110082001-0018</host>
  <port>24002</port>
</node>
<node index="3">
  <host>server-2110082001-0017</host>
  <port>24002</port>
</node>
</zookeeper>
```

- If yes, go to [Step 11](#).
- If no, go to [Step 10](#).

Step 10 Change the port number to the ZooKeeper port number, restart the ClickHouseServer instance, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

NOTICE

During the instance restart, the instance is unavailable, and the ClickHouse service on that instance fails to be executed.

Collect fault information.

Step 11 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 12 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 13 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 14 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 15 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.335 ALM-45441 Zookeeper Disconnected

Alarm Description

The system checks the connection between ClickHouse and ZooKeeper every minute. This alarm is generated when the connection fails. The alarm is reported because the ZooKeeper connection is abnormal. If the connection fails for three consecutive times, the system generates an alarm.

This alarm is automatically cleared when the system detects that the connection is normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45441	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If ClickHouse is disconnected from ZooKeeper, the ClickHouse service cannot be used.

Possible Causes

- The ZooKeeper service is abnormal.
- The ClickHouse service is overloaded.

Handling Procedure

Check whether ZooKeeper is normal.

- Step 1** On FusionInsight Manager, choose **Cluster > Services > ZooKeeper > quorumpeer**.
- Step 2** Check whether ZooKeeper instances are normal.
- If yes, go to [Step 6](#).
 - If no, go to [Step 3](#).
- Step 3** Select instances whose status is not good and choose **More > Restart Instance**.

NOTICE

During the instance restart, the instance is unavailable, and the ZooKeeper service on the current instance node fails to be executed.

- Step 4** Check whether the instance status is good after restart.
- If yes, go to [Step 5](#).
 - If no, go to [Step 10](#).
- Step 5** Choose **O&M > Alarm > Alarms** and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 6](#).

Check whether the ClickHouse service load is heavy.

- Step 6** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.
- Step 7** Log in to the node where the client is installed as the client installation user and run the following commands:

```
cd {Client installation path}
```

```
source bigdata_env
```

- For a cluster with Kerberos authentication enabled (security mode):
kinit *Component service user*
clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**

- For a cluster with Kerberos authentication disabled (normal mode):
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm **--user** *User name* **--password** **--port** 9440

Step 8 Run the following statement to check whether data is frequently written to the system table. If yes, wait until the service execution is complete and check whether the alarm is cleared.

```
SELECT query_id, user, FQDN(), elapsed, query FROM system.processes ORDER BY query_id;
```

- If yes, no further action is required.
- If no, go to [Step 9](#).

Step 9 Check whether a large amount of data is written. If yes, wait until the task is complete and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 12 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 13 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.336 ALM-45442 Too Many Concurrent SQL Statements

Alarm Description

The alarm module checks the number of concurrent ClickHouse requests every 30 seconds. This alarm is generated when the number of concurrent ClickHouse requests exceeds the concurrency threshold configured on the UI.

This alarm is cleared when the system detects that the actual number of concurrent requests is less than concurrency threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45442	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If there are too many concurrent SQL statements, a large number of system resources are consumed. As a result, system response becomes slow.

Possible Causes

The ClickHouse service is overloaded.

Handling Procedure

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.
- Step 2** Choose **Cluster > ClickHouse > Instance**, select an instance based on the alarm information. Choose **Chart > Concurrency** to check whether the actual number of concurrent SQL statements is greater than SQL concurrency threshold.
 - If yes, go to **Step 3**.
 - If no, go to **Step 5**.
- Step 3** Confirm with the user whether a large number of tasks were being executed during the alarming period.
 - If yes, go to **Step 4**.
 - If no, go to **Step 5**.
- Step 4** On FusionInsight Manager, choose **O&M** and click **Alarm > Thresholds** in the navigation pane on the left. On the displayed page, click **ClickHouse > Concurrency** and adjust the threshold, or wait until the task is complete. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

- Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 6** Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.
- Step 7** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.
- Step 8** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.337 ALM-45443 Slow SQL Queries in the Cluster

Alarm Description

The system checks slow SQL queries for ClickHouse every 1 minute. This alarm is generated when the execution time of a SQL statement is longer than or equal to the slow SQL threshold.

This alarm is automatically cleared when the system detects that the execution time of the SQL statement is shorter than the slow SQL threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45443	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The performance of the ClickHouse service deteriorates, which slows the response of other services. If there are too many slow SQL statements, the service may be unavailable.

Possible Causes

- The ClickHouse service is overloaded.
- The execution of SQL statements takes a long time.

Handling Procedure

Check whether the ClickHouse service load is heavy.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

```
cd {Client installation path}
```

```
source bigdata_env
```

- For a cluster with Kerberos authentication enabled (security mode):

```
kinit Component service user
```

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --port --secure
```

- For a cluster with Kerberos authentication disabled (normal mode):

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --user Username --password --port
```

Step 3 Run the following statement to check whether data is frequently written to the system table. If yes, wait until the service execution is complete and check whether the alarm is cleared.

```
SELECT query_id, user, FQDN(), elapsed, query FROM system.processes ORDER BY query_id;
```

- If yes, no further action is required.

- If no, go to [Step 4](#).

Checking whether the SQL statements take a long time.

Step 4 Check the logical cluster to which the alarm object belongs. Log in to FusionInsight Manager, click **Cluster**, choose **Services > ClickHouse**, and click **Logic Cluster**. On the displayed page, choose **Query Management > Ongoing Slow Queries**. Check which SQL statements take a long time on the displayed page, confirm with the user to adjust services, optimize slow SQL statements, and check whether the optimization is successful.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

Step 5 After the SQL statements are complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 7 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 8 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.338 ALM-45444 Abnormal ClickHouse Process

Alarm Description

The health check module checks ClickHouse instances every 30 seconds. If the number of consecutive failures exceeds the threshold, an alarm is reported. In this case, the ClickHouse process may stop responding and services cannot be properly executed.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45444	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the ClickHouse process is abnormal, services cannot run properly.

Possible Causes

The ClickHouse process runs improperly.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

```
cd {Client installation path}
```

```
source bigdata_env
```

- For a cluster with Kerberos authentication enabled (security mode):

```
kinit Component service user
```

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --port 9440 --secure
```

- For a cluster with Kerberos authentication disabled (normal mode):

```
clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm --user Username --password --port 9000
```

Step 3 Run the following statement to check whether the result can be properly returned:

SELECT 1;

- If yes, go to [Step 4](#).
- If no, go to [Step 5](#).

Step 4 Wait for several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 7 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.339 ALM-45585 IoTDB Service Unavailable

Alarm Description

The system checks the IoTDB service status every 300 seconds. This alarm is generated when the IoTDB service is unavailable. This alarm is cleared when the IoTDB service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45585	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The IoTDB service is unavailable for read and write operations.

Possible Causes

- The KrbServer service is abnormal.
- More than 50% of IoTDBServer instances are faulty.

Handling Procedure

Check whether the KrbServer service on which the IoTDB depends is abnormal.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 2 In the alarm list, check whether ALM-25500 KrbServer Service Unavailable exists.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Handle the alarm by referring to "ALM-25500 KrbServer Service Unavailable."

Step 4 After ALM-25500 is cleared, wait a few minutes and check whether the alarm HetuServer Service Unavailable is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check whether IoTDBServer instances are faulty.

Step 5 On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Service > IoTDB > Instance**.

Step 6 Check whether the percentage of faulty IoTDBServer instances exceeds 50%. If yes, restart the faulty IoTDBServer instances and check whether the status is restored.


- If yes, no further action is required.

- If no, go to [Step 7](#).

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 8 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.340 ALM-45586 IoTDBServer Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the IoTDBServer process status every 60 seconds. The alarm is generated when the heap memory usage of the IoTDBServer process exceeds the threshold (90% of the maximum memory).

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45586	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The read and write performance of the IoTDBServer process deteriorates, memory overflow occurs, and the IoTDBServer process breaks down.

Possible Causes

The heap memory of the IoTDBServer process is overused or the heap memory is inappropriately allocated.

Handling Procedure


Check the heap memory usage.

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **45586**, view the role name in **Location**, and check the instance IP address.
- Step 2** Choose **Cluster > Name of the desired cluster > Service > IoTDB > Instance**. Click the IoTDBServer for which the alarm is generated to go to **Dashboard**. Click the drop-down list in the upper right corner of the chart area and choose **Customize > Memory**. In the dialog box that is displayed, select **IoTDBServer Heap Memory Resource Percentage**, and click **OK**. Check whether the used non-heap memory of the IoTDBServer process reaches 90% (by default) of the maximum non-heap memory specified for IoTDBServer.
- If yes, go to **Step 3**.
 - If no, go to **Step 5**.
- Step 3** Choose **Cluster > Name of the desired cluster > Service > IoTDB > Configuration**, click **All Configurations**, choose **IoTDBServer > System**, and increase the value of **-Xmx** in the **GC_OPTS** parameter.

NOTE

- The default value of **-Xmx** is **2G**.
 - If this alarm is occasionally generated, increase the value by 0.5 times. If this alarm is frequently generated, double the value.
 - In the case of large service volume and high service concurrency, you are advised to add instances.
- Step 4** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to **Step 5**.

Collect the fault information.

- Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 6** Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.
- Step 7** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 8** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.341 ALM-45587 IoTDBServer GC Duration Exceeds the Threshold

Alarm Description

The system checks the GC duration of the IoTDBServer process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default) for three consecutive times. You can choose **O&M > Alarm > Threshold Configuration > Name of the desired cluster > IoTDB > GC > Total GC duration of IoTDBServer process (IoTDBServer)** to change the threshold. This alarm is cleared when the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45587	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The IoTDBServer process may be unavailable for read and write operations.

Possible Causes

The heap memory of the IoTDBServer process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **45587**, view the role name in **Location**, and check the instance IP address.
- Step 2** Choose **Cluster > Name of the desired cluster > Service > IoTDB > Instance**. Click the IoTDBServer for which the alarm is generated to go to **Dashboard**. Click the drop-down list in the upper right corner of the chart area and choose **Customize > GC**. In the dialog box that is displayed, select **Garbage Collection (GC) Time of IoTDBServer**, and click **OK**. Check whether the GC time of the IoTDBServer process is greater than 12 seconds.
- If yes, go to **Step 3**.
 - If no, go to **Step 5**.
- Step 3** Choose **Cluster > Name of the desired cluster > Service > IoTDB > Configuration**, click **All Configurations**, choose **IoTDBServer > System**, and increase the value of **-Xmx** in the **GC_OPTS** parameter.

NOTE


- The default value of **-Xmx** is **2G**.
 - If this alarm is occasionally generated, increase the value by 0.5 times. If this alarm is frequently generated, double the value.
 - In the case of large service volume and high service concurrency, you are advised to add instances.
- Step 4** Check whether the alarm is cleared.
- If yes, no further action is required.

- If no, go to [Step 5](#).

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.342 ALM-45588 IoTDBServer Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the IoTDBServer service every 60 seconds. This alarm is generated when the direct memory usage of the IoTDBServer instance exceeds the threshold (90% of the maximum memory) for five consecutive times. This alarm is cleared when the IoTDBServer direct memory usage is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45588	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause service unavailability, and the IoTDBServer process may be unavailable for read and write operations.

Possible Causes

The direct memory of the IoTDBServer process is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

- Step 1** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **45588**, view the role name in **Location**, and check the instance IP address.
- Step 2** Choose **Cluster > Name of the desired cluster > Service > IoTDB > Instance**. Click the IoTDBServer for which the alarm is generated to go to **Dashboard**. Click the drop-down list in the upper right corner of the chart area and choose **Customize > Memory**. In the dialog box that is displayed, select **IoTDBServer Direct Buffer Resource Percentage**, and click **OK**.
- Step 3** Check whether the direct memory used by the IoTDBServer reaches the threshold (90% of the maximum direct memory by default).
 - If yes, go to **Step 4**.
 - If no, go to **Step 6**.
- Step 4** On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Service > IoTDB > Configuration**, click **All Configurations**, choose **IoTDBServer > System**, increase the value of **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter as required, and save the configuration.

 NOTE

- If this alarm is generated, the direct memory configured for the IoTDBServer process cannot meet the requirements of the IoTDBServer process.
- You are advised to set **-XX:MaxDirectMemorySize** in **GC_OPTS** to twice the direct memory used by the IoTDBServer process. (You can change the value based on the actual service scenario.)
- To obtain the size of the direct memory used by the IoTDBServer process, choose **Customize > Memory > IoTDBServer Direct Memory Resource Status**. If **GC_OPTS** does not contain the **-XX:MaxDirectMemorySize** parameter, add it manually.


Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 7 Expand the **Service** drop-down list, and select **IoTDBServer** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.343 ALM-45589 ConfigNode Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the ConfigNode process every 60 seconds. This alarm is generated when the heap memory usage of the ConfigNode process exceeds the threshold (90% of the maximum memory). This alarm is cleared when the heap memory usage of the ConfigNode process is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45589	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System


The read and write performance of the ConfigNode process deteriorates, memory overflow occurs, and the ConfigNode process breaks down.

Possible Causes

The heap memory configured for the node is improper. As a result, the usage exceeds the threshold.

Handling Procedure

Check the heap memory configuration.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the real-time alarm list, click  in front of this alarm and view the role name and instance IP address in **Location**.
- Step 2** Choose **Cluster > Services > IoTDB**. Click **Instance**, click the ConfigNode corresponding to the IP address obtained in **Step 1**, and check whether **ConfigNode Heap Memory Usage** on the **Dashboard** tab page reaches the threshold specified for the ConfigNode process.

If the chart is not displayed, click the drop-down list in the upper right corner of the chart area and choose **Customize > Memory**. In the dialog box that is displayed, select **ConfigNode Heap Memory Usage** and click **OK**.

 NOTE

You can choose **O&M > Alarm > Threshold Configuration > *Name of the desired cluster* > IoTDB > Memory > ConfigNode Heap Memory Usage (ConfigNode)** to view the threshold.

- If yes, go to [Step 3](#).
- If no, go to [Step 6](#).

Step 3 Choose **Cluster > Services > IoTDB**. Click **Configurations** then **All Configurations**, click **ConfigNode**, and choose **System**. Set **-Xmx** in **GC_OPTS** to a larger value and save the configuration.

 NOTE

- The default value of **-Xmx** is **2G**.
- If this alarm is occasionally generated, increase the value of **-Xmx** by 0.5 times. If this alarm is frequently generated, double the value of **-Xmx**.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Click **Dashboard**. Click **Restart Service** to restart the IoTDB service for the configuration to take effect.

NOTICE

During the restart of the IoTDB service, read and write requests are interrupted.

Step 5 Wait for about 120 seconds and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M > Log > Download**.

Step 7 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 8 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.344 ALM-45590 ConfigNode GC Duration Exceeds the Threshold

Alarm Description

The system checks the GC duration of the ConfigNode process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default) for three consecutive times. This alarm is cleared when the GC duration is less than the threshold.

NOTE

You can choose **O&M > Alarm > Threshold Configuration > Name of the desired cluster > ioTDB > GC > Total GC duration of ConfigNode process (ConfigNode)** to increase the threshold by 20% each time.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45590	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System


The read and write performance of the ConfigNode process may deteriorate.

Possible Causes

The heap memory configured on the node is improper. As a result, GC occurs frequently.

Handling Procedure

Check the heap memory configuration.

Step 1 Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the real-time alarm list, click  in the row containing this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **Cluster > Services > IoTDB**. Click **Instance**, click the ConfigNode corresponding to the IP address obtained by [Step 1](#). Switch to the **Dashboard** tab page, locate the **Total GC Duration of ConfigNode** chart, and check whether the GC duration of the ConfigNode process exceeds the threshold.

If the GC duration of ConfigNode is not displayed, click the drop-down list in the upper right corner of the chart area and choose **Customize > GC**. In the displayed dialog box, select **Total GC Duration of ConfigNode** and click **OK**.

NOTE

You can choose **O&M > Alarm > Threshold Configuration > Name of the desired cluster > IoTDB > GC > Total GC duration of ConfigNode process (ConfigNode)** to view the threshold.

- If yes, go to [Step 3](#).
- If no, go to [Step 6](#).

Step 3 Choose **Cluster > Services > IoTDB**. Click **Configurations** then **All Configurations**, click **ConfigNode**, and choose **System**. Set **-Xmx** in **GC_OPTS** to a larger value and save the configuration.

NOTE

- The default value of **-Xmx** is **2G**.
- If this alarm is occasionally generated, increase the value by 0.5 times. If this alarm is frequently generated, double the value.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Click **Dashboard**. Click **Restart Service** to restart the IoTDB service for the configuration to take effect.


NOTICE

During the restart of the IoTDB service, read and write requests are interrupted.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 7** Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.
- Step 8** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.345 ALM-45591 ConfigNode Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the ConfigNode process every 60 seconds. This alarm is generated when the direct memory usage of the ConfigNode exceeds the threshold for five consecutive times. That is, the direct memory configured for ConfigNode cannot meet service requirements. This alarm is cleared when the direct memory usage of ConfigNode is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45591	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System


Direct memory overflow may cause instance unavailability, and the ConfigNode process may be unavailable for read and write operations.

Possible Causes

The direct memory configured for the node is improper. As a result, the usage exceeds the threshold.

Handling Procedure

Check the direct memory configuration.

Step 1 Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the real-time alarm list, click  in front of this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **Cluster > Services > IoTDB**. Click **Instance**, click the ConfigNode corresponding to the IP address obtained in **Step 1**, and check whether **ConfigNode Direct Memory Usage** on the **Dashboard** tab page reaches the threshold specified for the ConfigNode process (90% of the maximum direct memory by default).

If the chart is not displayed, click the drop-down list in the upper right corner of the chart area and choose **Customize > Memory**. In the dialog box that is displayed, select **ConfigNode Direct Memory Usage** and click **OK**.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > IoTDB**. Click **Configurations** then **All Configurations**. Click **ConfigNode** and select **System**. Set **-XX:MaxDirectMemorySize** in **GC_OPTS** to a larger value as required and save the configuration.

 NOTE

- You are advised to set **-XX:MaxDirectMemorySize** in **GC_OPTS** to twice the direct memory used by the ConfigNode process. (You can change the value based on the actual service scenario.)
- To obtain the size of the direct memory used by the ConfigNode process, choose **Customize > Memory > ConfigNode Direct Memory Resource Status**.
- If **GC_OPTS** does not contain the **-XX:MaxDirectMemorySize** parameter, add it.

Step 4 Restart the affected IoTDB service or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

NOTICE


During the restart of the IoTDB service or instance, read and write requests are interrupted.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M > Log > Download**.

Step 6 Expand the **Service** drop-down list, and select **ConfigNode** for the destination cluster.

Step 7 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.346 ALM-45592 IoTDBServer RPC Execution Duration Exceeds the Threshold

Alarm Description

The system checks the RPC execution duration of the IoTDBServer process every 60 seconds. This alarm is generated when the execution duration exceeds the

threshold. This alarm is cleared when the RPC execution time of the IoTDBServer process is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45592	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System


The read and write performance of the IoTDBServer process deteriorates.


Possible Causes

The processing duration of an IoTDBServer RPC request exceeds the threshold. Logs need to be further analyzed to locate the cause.

Handling Procedure

Collect fault information.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the real-time alarm list, click  in front of this alarm and view the role name and instance IP address in **Location**.
- Step 2** Choose **O&M > Log > Download**.
- Step 3** Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.
- Step 4** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in [Step 1](#), and click **OK**.

Step 5 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 6 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.347 ALM-45593 IoTDBServer Flush Execution Duration Exceeds the Threshold

Alarm Description

This alarm is generated when the data flush duration exceeds the threshold. This alarm is cleared when the flush duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45593	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System



The data writes are blocked, which causes write performance deterioration.

Possible Causes

The IoTDB flushing on the node is slow. You need to further analyze logs.

Handling Procedure

Collect fault information.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the real-time alarm list, click  in front of this alarm and view the role name and instance IP address in **Location**.
- Step 2** Choose **O&M > Log > Download**.
- Step 3** Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.
- Step 4** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in [Step 1](#), and click **OK**.
- Step 5** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 6** Contact O&M personnel and provide the collected logs.
- End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.348 ALM-45594 IoTDBServer Intra-Space Merge Duration Exceeds the Threshold

Alarm Description

This alarm is generated when the merge duration in the space exceeds the threshold. This alarm is cleared when the merge duration in the space is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45594	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System



Data write is blocked and the write operation performance is affected.

Possible Causes

The merge task in the IoTDB space of the node is slow. You need to further analyze logs.

Handling Procedure

Collect fault information.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the real-time alarm list, click  in front of this alarm and view the role name and instance IP address in **Location**.
- Step 2** Choose **O&M > Log > Download**.
- Step 3** Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.
- Step 4** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in [Step 1](#), and click **OK**.
- Step 5** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 6** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.349 ALM-45595 IoTDBServer Cross-Space Merge Duration Exceeds the Threshold

Alarm Description

This alarm is generated when the cross-space merge duration exceeds the threshold. This alarm is cleared when the cross-space merge duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45595	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System



Data writes are blocked, which causes write performance deterioration.

Possible Causes

The IoTDB cross-space merge task on the node is slow. You need to further analyze logs.

Handling Procedure

Collect fault information.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the real-time alarm list, click  in front of this alarm and view the role name and instance IP address in **Location**.
- Step 2** Choose **O&M > Log > Download**.
- Step 3** Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.
- Step 4** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in [Step 1](#), and click **OK**.
- Step 5** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 6** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.350 ALM-45596 Procedure Execution Failed

Alarm Description

Procedures are the tasks managed and executed by the ConfigNode leader. This alarm is generated when a procedure fails to be executed. This alarm is cleared when the procedure is successfully executed.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45596	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
ProcedureInformation	Specifies the procedure-related information.

Impact on the System

Operation consistency is affected.

Possible Causes

- The task for adding IoTDB replicas fails to be executed.
- The task for deleting the storage group fails to be executed.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**, locate this alarm, and click .

Step 2 Check the value of **ProcedureInformation** in **Location**. The value starts with the procedure type and contains main information about the procedure.

Check whether the task for adding replicas fails.

Step 3 Check whether the value of **ProcedureInformation** starts with **AddRegionProcedure** or **ReJoinDataNodeProcedure**.

- If yes, the task fails. Go to [Step 4](#).
- If no, go to [Step 5](#).

Step 4 Wait for half an hour. If the region is successfully added, the alarm is automatically cleared. Otherwise, go to [Step 5](#).

Check whether the task for deleting the storage group fails.


Step 5 Check whether the value of **ProcedureInformation** starts with **DeleteStorageGroupProcedure**.

- If yes, the storage group fails to be deleted. Go to [Step 6](#).
- If no, go to [Step 7](#).

Step 6 Delete the storage group displayed in **ProcedureInformation** again on the IoTDB client. If the deletion is successful, the alarm is automatically cleared. Otherwise, go to [Step 7](#).

Collect fault information.

Step 7 Choose **Cluster > Services > IoTDB > Instance** to view the hosts where all IoTDBServer and ConfigNode instances are located.

- Step 8** Choose **O&M > Log > Download**.
- Step 9** Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.
- Step 10** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in [Step 7](#), and click **OK**.
- Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 12** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.351 ALM-45615 CDL Service Unavailable

Alarm Description

The system checks the CDL health status every 60 seconds. This alarm is generated when the CDL health status is **DOWN**. This alarm is cleared when the system detects that the CDL health status is **UP**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45615	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The CDL service is abnormal. You cannot use FusionInsight Manager to perform cluster operations on the CDL service. The CDL service function is unavailable.

Possible Causes

All CDLService or CDLConnector instances of the CDL service are abnormal, and the Kafka service is unavailable.

Handling Procedure


Check whether the Kafka service on which the CDL service depends is abnormal.

- Step 1** On FusionInsight Manager, choose **O&M > Alarm > Alarms**.
- Step 2** In the alarm list, check whether ALM-38000 Kafka Service Unavailable exists.
- If yes, go to **Step 3**.
 - If no, go to **Step 5**.
- Step 3** Handle the alarm by referring to "ALM-38000 Kafka Service Unavailable".
- Step 4** After the alarm is cleared, wait a few minutes and check whether the alarm HetuServer Service Unavailable is cleared.
- If yes, no further action is required.
 - If no, go to **Step 5**.

Check whether CDL instances are faulty.

- Step 5** On FusionInsight Manager, choose **Cluster > Name of the desired cluster > Service > CDL > Instance**.
- Step 6** Check whether all CDLService and CDLConnector instances are faulty.
- If yes, restart the CDL service and choose **Cluster > Name of the desired cluster > Services > CDL > More > Restart Service**. If the fault persists after the restart, go to **Step 7** and contact O&M personnel to check CDL logs.
 - If no, go to **Step 7**.

Collect the fault information.

- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 8** Expand the **Service** drop-down list, and select **CDL** for the target cluster.
- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.
- Step 10** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the service is restored, the system automatically clears the alarm.

Related Information

None

7.12.352 ALM-45616 CDL Job Execution Exception

Alarm Description

The system checks whether a CDL job is normal every 60 seconds. This alarm is reported when the CDL job is abnormal. This alarm is cleared when the job is restored or stopped.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45616	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.
Username	Specifies the username of the job for which the alarm was generated.


Impact on the System


CDL tasks fail, and real-time data integration is interrupted.

Possible Causes

The CDL task fails to be executed due to incorrect parameter settings or other reasons. On the **Job Management** page of the CDL web UI, locate the row where the job is located and click **Failed/Abnormal running** in the **Status** column to view the failure cause, or view the failure cause in the logs.

Handling Procedure

- Step 1** Log in to FusionInsight Manager as a user who has the CDL job creation or administrator permission.
- Step 2** Choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**, click  in the row where **Alarm ID** is **45616**, and view the name of the job for which this alarm is generated in **Location**.
- Step 3** Choose **Cluster > Services > CDL** and click the link next to **CDLService UI** to go to the CDL web UI.
- Step 4** Locate the row where the failed job is located based on the job name obtained in **Step 2**, and click **Abnormal running** or **Failed** in the **Status** column.

Name	Created	Status	Type
pghudi		 Abnormal running	pgsql ----> kafka ----> hudi

- Step 5** On the page that is displayed, view the error information and rectify the fault. For example, **Figure 7-180** shows that the task running on Yarn is manually killed. For details, see trace error information, as shown in **Figure 7-181**.

Figure 7-180 CDL job exception

Task Details

Basic Information

job-name		submission-id	5	execution-start-time	2022-01-11 14:15
app-id	application_1640579034647_0077	app-status	KILLED		

Source information

source-connector-id	3	source-connector-name	pghudi---3---5
---------------------	---	-----------------------	----------------

type	work.id	task.id	state	trace
connector		NA	RUNNING	
task		0	RUNNING	

Sink information

sink-connector-id

Figure 7-181 Trace error information

Task Details

Basic Information

job-name		submission-id	231	execution-start-time	
----------	--	---------------	-----	----------------------	--

Source information

source-connector-id	99	source-connector-name	
---------------------	----	-----------------------	--

type	work.id	task.id	state	trace
connector		NA	RUNNING	
task		0	FAILED	java.lang.RuntimeException: org.apache.kafka.connect.errors.Con...

Sink information


sink-connector-id

- Step 6** Rectify the fault based on the error information, execute the task again, and check whether the task can be executed successfully.
- If yes, no further action is required.
 - If no, go to [Step 7](#).

Collect the fault information.

- Step 7** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 8** Select **CDL** in the required cluster for **Service**.

- Step 9** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

- Step 10** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the job is successfully restored or stopped, the alarm is cleared if it has been reported.

Related Information

None

7.12.353 ALM-45617 Data Queued in the CDL Replication Slot Exceeds the Threshold

Alarm Description

If too many WALs are stacked in PostgreSQL or openGauss (applicable to MRS 3.3.0 or later), the PostgreSQL or openGauss disk space may be used up. The system checks whether the amount of data queued in the replication slot configured for a CDL job exceeds the threshold every 5 minutes. This alarm is generated when the amount of data queued in the replication slot exceeds the threshold. This alarm is cleared when the number of data queued in the replication slot falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45617	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.
DBName	Specifies the database for which the alarm was generated.
SlotName	Specifies the database replication slot for which the alarm was generated.
Lag	Specifies the data queued in the slot.


Impact on the System

Disk space of the source PostgreSQL or openGauss database may be used up and the database cannot provide services.

Possible Causes

The CDL job is abnormal, and data processing stops; the source database is updated quickly, and CDL data processing is slow.

Handling Procedure

- Step 1** Log in to FusionInsight Manager as a user who has the CDL job creation or administrator permission.
- Step 2** Choose **O&M**. In the navigation pane on the left, choose **Alarm > Alarms**, click  in the row where **Alarm ID** is **45617**, and view the name of the job for which this alarm is generated in **Location**.
- Step 3** Check whether **ALM-45616 CDL Job Execution Exception** is displayed in the alarm list.
 - If yes, handle the alarm by performing operations provided for **ALM-45616 CDL Job Execution Exception**.
 - If no, go to **Step 4**.
- Step 4** Choose **Cluster > Services > CDL**. Click the link next to **CDLService UI** to go to the CDL web UI and check whether the job is displayed in the job list based on its name obtained in **Step 2**.
 - If yes, check whether the job is abnormal.
 - If it is abnormal, go to **Step 5**.
 - If it is not, data processing is slow. Contact O&M personnel.

- If no, go to [Step 7](#).

Step 5 Click **Abnormal** or **Failed** in the row where the job is located and rectify the fault based on the error information displayed on the page.


Step 6 After rectifying the fault, run the job again and check whether the job can be executed successfully.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 8 Expand the **Service** drop-down list, and select **CDL** for the target cluster.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is cleared when the amount of data queued in the replication slot is less than the threshold. You do not need to manually clear the alarm.

Related Information

None

7.12.354 ALM-45635 FlinkServer Job Execution Failure

This section applies to MRS 3.1.2 or later.

Alarm Description

The system checks whether FlinkServer jobs fail to be executed every 10 seconds. This alarm is generated when a FlinkServer job fails. This alarm is cleared when the job is successfully restarted.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45635	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.

Impact on the System

This alarm is a job-level alarm and does not affect FlinkServer. You need to view Flink job logs to find out the failure cause.


Possible Causes

You can view failure causes in specific logs.

Handling Procedure

- Step 1** Log in to Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.
- Step 3** Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 7-182 Application ID of a job



If yes, go to [Step 4](#).

If no, go to [Step 6](#).

- Step 4** Click the application ID of the failed job to go to the job page.
 1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-183 Clicking Logs

Show 20 entries					
Attempt ID	Started	Node	Logs		
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0	

Showing 1 to 1 of 1 entries

- Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 7-184 Clicking the ID in the Attempt ID column

Show 20 entries					
Attempt ID	Started	Node	Logs		
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0	

Showing 1 to 1 of 1 entries

Figure 7-185 Clicking Logs

Show 20 entries					
Container ID	Node	Container Exit Status	Logs		
container_0009_01_000002	https/	0	Logs		
container_0009_01_000001	https/	0	Logs		

Showing 1 to 2 of 2 entries

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster > Services > Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More > Job Monitoring** in the **Operation** column to view the TaskManager logs.

- Step 5** View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

- Step 6** On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities > Browse the file system**, and download logs in the `/tmp/logs/User name/logs/Application ID of the failed job` directory.

- Step 7** View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearance

After the job is successfully restarted, the alarm is cleared if it has been reported.

Related Information

None

7.12.355 ALM-45636 Flink Job Checkpoints Keep Failing

This section applies to MRS 3.1.2 or a version between 3.1.2 and 3.3.0.

Description

The system checks the number of consecutive checkpoint failures based on the configured alarm checking interval. This alarm is generated when the number of consecutive checkpoint failures of a FlinkServer job reaches the configured threshold. This alarm is cleared when checkpoints are recovered or the job is successfully restarted.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45636	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.
Username	Specifies the username of the job for which the alarm is generated.

Impact on the System

The Flink job may fail. You need to check the status and logs of the Flink job to locate the fault. This is a job-level alarm and has no impact on FlinkServer.

Possible Causes

You can view failure causes in specific logs.

Procedure

- Step 1** Log in to Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.

- Step 3** Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 7-186 Application ID of a job

ID	User	Name	Application Type	Queue
application_1_0009	f...	zw_..._kafka	Apache Flink	default

If yes, go to **Step 4**.

If no, go to **Step 6**.

- Step 4** Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-187 Clicking Logs

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 7-188 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

Figure 7-189 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https://-	0	Logs
container_0009_01_000001	https://-	0	Logs

Showing 1 to 2 of 2 entries

 NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster > Services > Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More > Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities > Browse the file system**, and download logs in the `/tmp/logs/User name/logs/Application ID of the failed job` directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

This alarm is cleared when Flink job checkpoints are recovered or the job is successfully restarted.

Related Information

None

7.12.356 ALM-45636 Number of Consecutive Checkpoint Failures of a Flink Job Exceeds the Threshold

This section applies to MRS 3.3.1 or later.

Alarm Description

The system checks the number of consecutive checkpoint failures based on the configured alarm checking interval. This alarm is generated when the number of consecutive checkpoint failures of a FlinkServer job reaches the configured threshold. This alarm is cleared when checkpoints are recovered or the job is successfully restarted.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45636	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	ApplicationName	Specifies the name of the application for which the alarm was generated.
	JobName	Specifies the job for which the alarm was generated.
	UserName	Specifies the username for which the alarm was generated.
Additional Information	ThresholdValue	Specifies the threshold value for triggering the alarm.
	CurrentValue	Specifies the value that triggered the alarm.

Impact on the System

The Flink job may fail. You need to check the status and logs of the Flink job to locate the fault. This is a job-level alarm and has no impact on FlinkServer.

Possible Causes

You can view failure causes in specific logs.

Handling Procedure

- Step 1** Log in to Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.
- Step 3** Locate the failed task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the native Yarn page.

Figure 7-190 Application ID of a job



- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-191 Clicking Logs

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 7-192 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

Figure 7-193 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https://-	0	Logs
container_0009_01_000001	https://-	0	Logs

Showing 1 to 2 of 2 entries

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M engineers and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the `/tmp/logs/Username/logs/Application ID of the failed job` directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M engineers and send the collected fault logs.

----End

Alarm Clearance

This alarm is cleared when FlinkServer job checkpoints are recovered or the job is successfully restarted.

Related Information

None.

7.12.357 ALM-45637 FlinkServer Task Is Continuously Under Back Pressure

This section applies to MRS 3.1.2 or later.

Alarm Description

The system checks the back pressure duration of FlinkServer tasks based on the configured alarm checking interval. This alarm is generated when the back pressure duration of a FlinkServer task reaches the configured threshold. This alarm is cleared when the task back pressure is recovered or the job is successfully restarted.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45637	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.

Impact on the System

Continuous back pressure of Flink jobs may cause performance problems or checkpoint failures. Flink jobs fail. You need to check the status and logs of the Flink jobs to locate the cause. This is a job-level alarm and has no impact on FlinkServer.

Possible Causes

You can view the causes in the specific logs.

Handling Procedure

- Step 1** Log in to Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.
- Step 3** Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 7-194 Application ID of a job

ID	User	Name	Application Type	Queue
application_1_0009	f...	zw_..._kafka	Apache Flink	default

If yes, go to **Step 4**.

If no, go to **Step 6**.

- Step 4** Click the application ID of the failed job to go to the job page.
 1. Click **Logs** in the **Logs** column to view JobManager logs.


Figure 7-195 Clicking Logs

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 7-196 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

Figure 7-197 Clicking Logs**NOTE**

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster > Services > Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More > Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities > Browse the file system**, and download logs in the */tmp/logs/User name/logs/Application ID of the failed job* directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is cleared when FlinkServer task back pressure is recovered or the job is successfully restarted.

Related Information

None

7.12.358 ALM-45638 Number of Restarts After FlinkServer Job Failures Exceeds the Threshold

This section applies to MRS 3.1.2 or a version between 3.1.2 and 3.2.0.

Alarm Description

The system checks the number of FlinkServer job restarts based on the alarm checking interval. This alarm is generated when the number exceeds the configured threshold. This alarm is cleared when the job is restarted.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45638	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.

Impact on the System

Flink jobs are frequently restarted due to the failures. You need to locate the cause. This is a job-level alarm and has no impact on FlinkServer.

Possible Causes

You can view the causes in the specific logs.

Handling Procedure

- Step 1** Log in to Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.
- Step 3** Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 7-198 Application ID of a job

ID	User	Name	Application Type	Queue
application_1	I_0009	zw_..._kafka	Apache Flink	default

If yes, go to [Step 4](#).

If no, go to [Step 6](#).

- Step 4** Click the application ID of the failed job to go to the job page.
 1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-199 Clicking Logs

Show 20 entries					
Attempt ID	Started	Node	Logs		
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0	

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 7-200 Clicking the ID in the Attempt ID column

Show 20 entries					
Attempt ID	Started	Node	Logs		
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0	

Showing 1 to 1 of 1 entries

Figure 7-201 Clicking Logs

Show 20 entries					
Container ID	Node	Container Exit Status	Logs		
container_0009_01_000002	https/	0	Logs		
container_0009_01_000001	https/	0	Logs		

Showing 1 to 2 of 2 entries

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster > Services > Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More > Job Monitoring** in the **Operation** column to view the TaskManager logs.

- Step 5** View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

- Step 6** On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities > Browse the file system**, and download logs in the `/tmp/logs/User name/logs/Application ID of the failed job` directory.

- Step 7** View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is cleared when the FlinkServer job is successfully restarted.

Related Information

None

7.12.359 ALM-45638 Number of Restarts After Flink Job Failures Exceeds the Threshold

This section applies to MRS 3.2.0 or later.

Description

The system checks the number of Flink job restarts based on the alarm checking interval. This alarm is generated when the number exceeds the configured threshold. This alarm is cleared when the job is restarted.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45638	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.
Username	Specifies the username of the job for which the alarm is generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

Flink jobs are frequently restarted due to the failures. You need to locate the cause. This is a job-level alarm and has no impact on FlinkServer.

Procedure

- Step 1** Log in to Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.

- Step 3** Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 7-202 Application ID of a job

ID	User	Name	Application Type	Queue
application_1_0009	f...	zw_..._kafka	Apache Flink	default

If yes, go to **Step 4**.

If no, go to **Step 6**.

- Step 4** Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-203 Clicking Logs

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 7-204 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

Figure 7-205 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https://-	0	Logs
container_0009_01_000001	https://-	0	Logs

Showing 1 to 2 of 2 entries

 NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster > Services > Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More > Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities > Browse the file system**, and download logs in the **/tmp/logs/User name/logs/Application ID of the failed job** directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs.

----End

Alarm Clearing

This alarm is cleared when the Flink job is successfully restarted.

Related Information

None

7.12.360 ALM-45639 Checkpointing of a Flink Job Times Out

Description

The system checks the checkpointing timeout of Flink jobs every 30 seconds. This alarm is generated if the checkpointing timeout of a Flink job is longer than the threshold (600 seconds by default). This alarm is cleared when the checkpointing timeout of a job is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45639	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName (available in MRS 3.2.1 or later)	Specifies the name of the application for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.
UserName	Specifies the username for which the alarm is generated.

Impact on the System

The checkpointing fails. You need to locate the cause. This is a job-level alarm and has no impact on FlinkServer.

Possible Causes

The job may be in the sub-healthy state. The possible causes are as follows:

- The memory for the TaskManager of the job is insufficient.
- The state memory is too large, making checkpointing time-consuming.

Procedure

- Step 1** Log in to Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **O&M > Alarm > Alarms > ALM-45639 Checkpointing of a Flink Job Times Out**, view **Location**, and obtain the name of the task for which the alarm is generated.
- Step 3** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.
- Step 4** Locate the failed task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-206 Application ID of a job



ID	User	QueueUser	Name
application			

- If yes, go to **Step 5**.

- If no, go to [Step 7](#).

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-207 Clicking Logs

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 7-208 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

Showing 1 to 1 of 1 entries

Figure 7-209 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https://-	0	Logs
container_0009_01_000001	https://-	0	Logs

Showing 1 to 2 of 2 entries

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 6 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 7 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the `/tmp/logs/Username/logs/Application ID of the failed job` directory.

Step 8 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

This alarm is cleared when the checkpointing timeout a Flink job is less than or equal to the threshold.

Related Information

None

7.12.361 ALM-45640 FlinkServer Heartbeat Interruption Between the Active and Standby Nodes

This section applies to MRS 3.2.0 or later.

Alarm Description

This alarm is generated when the FlinkServer active node or standby node does not receive heartbeat messages from the peer for 30 seconds (heartbeat interruption duration configured in keepalive).

This alarm is cleared when the heartbeat recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45640	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The impact varies depending on the cause. If the heartbeat is interrupted due to other reasons, for example, network problems, two active nodes may exist because the standby node became the active node. Data synchronization between the active and standby nodes is abnormal, but FlinkServer can still provide services.

Possible Causes

- The active or standby FlinkServer instance is in the stopped state.
- The NIC of the floating IP address of the HA system used by the FlinkServer node is incorrectly configured. FlinkServer fails to be started.
- The link between the active and standby FlinkServer nodes is abnormal.

Handling Procedure

Check the status of the active and standby FlinkServer instances.

- Step 1** Log in to FusionInsight Manager, choose **Cluster > Services > Flink > Instance**, and check the state of FlinkServer is normal.
- If yes, go to [Step 3](#).
 - If no, go to [Step 2](#).
- Step 2** Select the abnormal FlinkServer instance and start the instance. After the instance is started, check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 3](#).

NOTICE

During the restart, the FlinkServer instance cannot provide services, but submitted jobs are not affected.

Check whether the link between the standby FlinkServer nodes is normal.

- Step 3** Choose **Cluster > Services > Flink > Instance**, and check the two service IP addresses of FlinkServer.
- Step 4** Log in to the server where the abnormal FlinkServer instance locates as the **root** user.
- Step 5** Run the following command to check whether the server of the other FlinkServer instance is reachable:
- ping** *IP address of the other FlinkServer instance*
- If yes, go to [Step 8](#).
 - If no, go to [Step 6](#).
- Step 6** Ask the network administrator to handle the network exception.
- Step 7** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 8](#).

Check whether the logs of the node where the abnormal FlinkServer instance locates contains error information.

- Step 8** Log in to the server where the abnormal FlinkServer instance locates as the **root** user.


- Step 9** Open the log file in the default directory `/var/log/Bigdata/flink/flinkserver/prestart.log` and check whether there is error message `Float ip x.x.x.x is invalid`.
- If yes, go to [Step 10](#).
 - If no, go to [Step 12](#).
- Step 10** On FusionInsight Manager, choose **Cluster > Services > Flink > Configurations > All Configurations** and search for `flink.ha.floatip`. Change the parameter value to the correct floating IP address, save the configuration, and restart the Flink service.

NOTICE

- Contact the network engineer to obtain the new floating IP address.
 - During the service restart, FlinkServer cannot provide services, but submitted jobs are not affected.
 - During the restart, the FlinkServer instance cannot provide services, but submitted jobs are not affected.
-

- Step 11** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 12](#).

Collect the fault information.

- Step 12** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 13** Select the Flink service in the required cluster for **Service**.
- Step 14** Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.
- Step 15** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 16** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.362 ALM-45641 Data Synchronization Exception Between the Active and Standby FlinkServer Nodes

This section applies to MRS 3.2.0 or later.

Alarm Description

The system checks data synchronization between the active and standby FlinkServer nodes every 60 seconds. This alarm is generated when the standby FlinkServer node fails to synchronize files with the active FlinkServer node.

This alarm is cleared when the standby FlinkServer synchronizes files with the active FlinkServer.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45641	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

After an active/standby switchover, some configurations may be lost. Some jobs and connections of the FlinkServer are interrupted, but the FlinkServer can still provide services properly.

Possible Causes

- The link between the active and standby FlinkServer nodes is interrupted.
- The synchronization file does not exist or the file permission is required.

Handling Procedure

Check whether the network between the active and standby FlinkServer is in normal state.

Step 1 On FusionInsight Manager, choose **Cluster > Services > ClickHouse > Instance**. View and record the IP addresses of active and standby FlinkServer.

Step 2 Log in to the active FlinkServer node as the **root** user.

Step 3 Run the following command to check whether the standby FlinkServer is reachable:

```
ping IP address of the standby FlinkServer
```

- If yes, go to [Step 6](#).
- If no, go to [Step 4](#).

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Check whether the storage space of the /srv/BigData/LocalBackup directory is insufficient.

Step 6 Run the following command to check whether the storage space of the /srv/BigData/LocalBackup directory is insufficient:

```
df -hl /srv/BigData/LocalBackup
```

- If yes, go to [Step 7](#).
- If no, go to [Step 10](#).

Step 7 Run the following command to clear unnecessary backup files:

```
rm -rf Directory to be cleared
```

The following are two examples:

```
rm -rf /srv/BigData/LocalBackup/0/default-oms_20191211143443
```

Step 8 On FusionInsight Manager, choose **O&M > Backup and Restoration > Backup Management**.

In the **Operation** column of the backup task, click **Configure** and change the value of **Maximum Number of Backup Copies** to reduce the number of backup file sets.

Step 9 Wait for 1 minute and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 10](#).

Check whether the synchronization file exists and whether the file permission is valid.

Step 10 Run the following command to check whether the synchronization file exists:

```
find /srv/BigData/ -name "sed*"
```

```
find /opt -name "sed*"
```

- If yes, go to [Step 11](#).

- If no, go to [Step 12](#).

Step 11 Run the following command to check the synchronization file information and permission queried in [Step 10](#):

ll *Path of the file you want to search for*

- If the file size is 0 and all values in the permission column are -, the file is a junk file. Run the following command to delete it:

```
rm -rf Files to be deleted
```

Wait for several minutes and check whether the alarm is cleared. If the alarm persists, go to [Step 12](#).

- If the file size is not 0, go to [Step 12](#).

Step 12 View the log file generated when the alarm is reported.

1. Run the following command to go to the HA run log file path of the current cluster:

```
cd /var/log/Bigdata/flink/flinkserver/ha/runlog
```

2. Decompress log file and view the logs generated when the alarm is reported.

For example, if the name of the file is **ha.log.2021-03-22_12-00-07.gz**, run the following command:

```
gunzip ha.log.2021-03-22_12-00-07.gz
```

```
vi ha.log.2021-03-22_12-00-07
```

Check whether error information is displayed before and after the alarm generation time in the logs.

- If it is displayed, rectify the fault based on the error information. Go to [Step 13](#).

For example, if the following error information is displayed, the directory permission is required. In this case, obtain the directory permission that is the same as the permission on a normal node.

```
2021-03-22 14:08:35.339 [10195489349] [0] INFO [add task(null) to list successful][HA][sync_module.c: SYNC_ActiveTask,1151][ha.bin,26572,35]
2021-03-22 14:08:35.339 [10195489349] [0] INFO [Start Task All Sync][HA][sync_core_inf.c:SYNC_StartTask,183][ha.bin,26572,35]
2021-03-22 14:08:35.339 [10195489349] [0] NOTICE [send sync task(alltask) to component successful][HA][sync_module.c: SYNC_SendsyncTask,832][ha.bin,26572,35]
2021-03-22 14:08:35.344 [10195489353] [0] INFO [open lstat failed:/opt/bigdata/apache-tomcat-7.0.70/conf/security/tomcat_on.crt ]. Permission denied.[HA]
gt.c: create_TravelName_Open,482][ha.bin,26572,41]
2021-03-22 14:08:35.344 [10195489353] [0] ERROR [Travel stack failed][HA][sync_fi[empt.c: Create_TravelName,613][ha.bin,26572,41]
2021-03-22 14:08:35.344 [10195489353] [0] ERROR [Increate list fail][HA][sync_filemt.c: SYNC_CreateFileList,285][ha.bin,26572,41]
2021-03-22 14:08:35.344 [10195489353] [0] ERROR [createFileList failed][HA][sync_core.c: SYNC_Task_SendEnd,1866][ha.bin,26572,41]
2021-03-22 14:08:35.344 [10195489353] [0] ERROR [[41][SendEnd][Task]Failed][HA][sync_core.c: SYNC_DbgMsgErr,202][ha.bin,26572,41]
2021-03-22 14:08:35.344 [10195489353] [0] ERROR [TaskEnd failed][HA][sync_core.c: SYNC_Err_TaskEnd,2723][ha.bin,26572,41]
2021-03-22 14:08:35.344 [10195489353] [0] NOTICE [hasendAlarm info: tcl1,category=0,cause=0,location=1,addinfo=1,location=1,location=(node-master1onFC) (locha=(192-168-
```

- If no, go to [Step 14](#).

Step 13 Wait for about 10 minutes and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 14](#).

Collect fault information.

Step 14 On FusionInsight Manager, choose **O&M > Log > Download**.

Step 15 Select FlinkServer information from **Services** and click **OK**.

Step 16 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 17 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.363 ALM-45642 RocksDB Continuously Triggers Write Traffic Limiting

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when RocksDB for a job continuously triggers write traffic limiting, that is, the RocksDB write rate is not 0. This alarm is cleared when the RocksDB write rate of the job becomes 0.

The **rocksdb.actual-delayed-write-rate** parameter specifies the RocksDB write rate of a job. Value **0** indicates that the rate is not limited, and other values indicate traffic limiting.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45642	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes

When the rate at which Flink jobs write data to RocksDB is not 0, write traffic limiting is triggered. The possible causes are as follows:

- There are too many MemTables. As a result, write traffic is limited or write stops, and **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** is generated.
- The size of SST files at level 0 is too large, and **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** is generated.
- The estimated compaction size exceeds the threshold, and **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** is generated.

Handling Procedure

Check whether write traffic limiting or write stop is caused due to too many MemTables.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 2 In the alarm list, check whether **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Handle the alarm by following the instructions provided in section **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45643 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check whether write traffic limiting or write stop is caused due to too many SST files at level 0.

Step 5 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 6 In the alarm list, check whether **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

Step 7 Handle the alarm by following the instructions provided in section **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**.

Step 8 After ALM-45644 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Check whether write traffic limiting or write stop is caused because the estimated compaction size exceeds the threshold.

Step 9 In the alarm list, check whether **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 10](#).
- If no, go to [Step 12](#).

Step 10 Handle the alarm by following the instructions provided in section **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 11 After ALM-45647 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

Collect fault information.

Step 12 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 13 Choose **O&M > Alarm > Alarms > ALM-45642 RocksDB Continuously Triggers Write Traffic Limiting**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 14 Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 15 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the native Yarn page.

Figure 7-210 Application ID of a job

Show 20 entries			
ID	User	QueueUser	Name
application			

- If yes, go to **Step 16**.
- If no, go to **Step 18**.

Step 16 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-211 Clicking Logs

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs.

Figure 7-212 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

Figure 7-213 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https://-	0	Logs
container_0009_01_000001	https://-	0	Logs

Showing 1 to 2 of 2 entries

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 17 View the job logs to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 18 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the */tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task* directory.

Step 19 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.364 ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the MemTable size of RocksDB for a job continuously exceeds the threshold (**metrics.reporter.alarm.job.alarm.rocksdb.get.micros.threshold**, 50000 microseconds by default). This alarm is cleared when the MemTable size of RocksDB for the job is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45643	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes

The write pressure of RocksDB is high.

Handling Procedure

Check TaskManager logs for the write pressure of RocksDB and collect logs.

- Step 1** Log in to FusionInsight Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **O&M > Alarm > Alarms > ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.
- Step 3** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.
- Step 4** Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-214 Application ID of a job

ID	User	QueueUser	Name
application_			

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-215 Clicking Logs

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 7**.

Figure 7-216 Clicking the ID in the Attempt ID column





Figure 7-217 Clicking Logs



NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the `/tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task` directory.

Check whether the write pressure of RocksDB is high.

Step 7 Check whether the value of `rocksdb.size-all-mem-tables` (unit: byte) in the TaskManager monitoring logs (keyword `RocksDBMetricPrint`) is greater than or equal to the total write buffer size (Total write buffer = `write_buffer_size` x `max_write_buffer_number`).

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to [Step 8](#).

Table 7-106 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.writebuffer.count	<ul style="list-style-type: none"> 2 4: enables SPINNING_DISK_OPTIMIZED_HIGH_MEM. 	<ul style="list-style-type: none"> Number of buffers 2 to 10 are recommended. Adjust the value based on service requirements.
state.backend.rocksdb.writebuffer.size	64MB	<ul style="list-style-type: none"> Buffer size 64MB to 256MB are recommended.

Parameter	Default Value	Description
state.backend.rocksdb.thread.num	<ul style="list-style-type: none">- 2- 4: enables SPINNING_DISK_OPTIMIZED_HIGH_MEM.	<ul style="list-style-type: none">- Number of flush threads. Increase the number of threads to quickly flush memory data to disks.- When the number of threads is increased, the number of vCores also needs to be increased.- 2 to 10 are recommended.

- If no, go to [Step 9](#).

Step 8 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Step 9 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.365 ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the number of SST files at level 0 of RocksDB for a job continuously exceeds the threshold (**state.backend.rocksdb.level0_slowdown_writes_trigger**, 20 by default). This alarm is cleared when the number of SST files at level 0 of RocksDB for the job is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45644	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes

Possible causes are as follows:

- The compaction pressure of RocksDB is too high, and **ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** and **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** are generated.
- There are too many SST files at level 0.

Handling Procedure

Check whether the compaction pressure of RocksDB is too high and ALM-45646 is generated.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 2 In the alarm list, check whether **ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 3](#).

- If no, go to [Step 5](#).

Step 3 Handle the alarm by following the instructions provided in section **ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45646 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check whether the compaction pressure of RocksDB is too high and ALM-45647 is generated.

Step 5 In the alarm list, check whether **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 6](#).
- If no, go to [Step 8](#).

Step 6 Handle the alarm by following the instructions provided in section **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 7 After ALM-45647 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

Check TaskManager logs for the number of SST files at level 0 and collect logs.

Step 8 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 9 Choose **O&M > Alarm > Alarms > ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 10 Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 11 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-218 Application ID of a job

ID	User	QueueUser	Name
application			

- If yes, go to [Step 12](#).
- If no, go to [Step 13](#).

Step 12 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-219 Clicking Logs

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to [Step 14](#).

Figure 7-220 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

Figure 7-221 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https://-	0	Logs
container_0009_01_000001	https://-	0	Logs

Showing 1 to 2 of 2 entries

NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 13 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the `/tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task` directory.

Check whether the number of SST files at level 0 is too large.

Step 14 Check whether the value of `rocksdb.num-files-at-level0` in TaskManager monitoring logs (keyword `RocksDBMetricPrint`) is greater than or equal to the value of `state.backend.rocksdb.level0_slowdown_writes_trigger` or `state.backend.rocksdb.level0_stop_writes_trigger`.

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to [Step 15](#).

Table 7-107 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.level0_slowdown_writes_trigger	20	<ul style="list-style-type: none">– Number of files that trigger slowdown at level 0– 20 to 30 are recommended.
state.backend.rocksdb.level0_stop_writes_trigger	36	<ul style="list-style-type: none">– Maximum number of files that trigger stop at level 0– 36 to 46 are recommended.

- If no, go to [Step 16](#).

Step 15 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 16](#).

Step 16 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.366 ALM-45645 Pending Flush Size of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the number of pending flush requests of RocksDB for a job continuously reaches n times the number of flush/compaction threads. This alarm is cleared when the number of pending flush requests of RocksDB for the job is less than or equal to the threshold.

- The number of flush/compaction threads is the value of **state.backend.rocksdb.thread.num**. The default value is **2**. If **SPINNING_DISK_OPTIMIZED_HIGH_MEM** is enabled, the default value is **4**.
- The **metrics.reporter.alarm.job.alarm.rocksdb.background.jobs.multiplier** parameter specifies n times the number of flush/compaction threads. The default value is **2**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45645	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes

The number of pending flush requests of RocksDB for the Flink job is too large.

Handling Procedure

Check TaskManager logs for the number of pending flush requests and collect logs.

- Step 1** Log in to FusionInsight Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **O&M > Alarm > Alarms > ALM-45645 Pending Flush Size of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.
- Step 3** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.
- Step 4** Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-222 Application ID of a job

ID	User	QueueUser	Name
application			

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-223 Clicking Logs

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 7**.

Figure 7-224 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

Figure 7-225 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https://-	0	Logs
container_0009_01_000001	https://-	0	Logs

NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and

download logs in the `/tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task` directory.

Check whether there are too many pending flush requests.

Step 7 Check whether the sum of the values of `rocksdb.mem-table-flush-pending` and `rocksdb.compaction-pending` in TaskManager monitoring logs (keyword `RocksDBMetricPrint`) is greater than n times the number of RocksDB threads (`metrics.reporter.alarm.job.alarm.rocksdb.background.jobs.multiplier`, 2 by default). If it is, you can increase the number of RocksDB threads.

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to [Step 8](#).

Table 7-108 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.thread.num	<ul style="list-style-type: none"> - 2 - 4: enables <code>SPINNING_DISK_OPTIMIZE_HIGH_MEMORY</code>. 	<ul style="list-style-type: none"> - Number of flush threads. Increase the number of threads to quickly flush memory data to disks. - When the number of threads is increased, the number of vCores also needs to be increased. - 2 to 10 are recommended.

- If no, go to [Step 9](#).

Step 8 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Step 9 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.367 ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the number of pending compaction requests of RocksDB for a job continuously reaches n times the number of flush/compaction threads. This alarm is cleared when the number of pending compaction requests of RocksDB for the job is less than or equal to the threshold.

- The number of flush/compaction threads is the value of **state.backend.rocksdb.thread.num**. The default value is **2**. If **SPINNING_DISK_OPTIMIZED_HIGH_MEM** is enabled, the default value is **4**.
- The **metrics.reporter.alarm.job.alarm.rocksdb.background.jobs.multiplier** parameter specifies n times the number of flush/compaction threads. The default value is **2**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45646	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes

The number of pending compaction requests of RocksDB for the Flink job is too large.

Handling Procedure

Check TaskManager logs for the number of pending compaction requests and collect logs.

- Step 1** Log in to FusionInsight Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **O&M > Alarm > Alarms > ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.
- Step 3** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.
- Step 4** Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-226 Application ID of a job

ID	User	QueueUser	Name
application			

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

- Step 5** Click the application ID of the failed job to go to the job page.
 1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-227 Clicking Logs


Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 7**.

Figure 7-228 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs

Figure 7-229 Clicking Logs



NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the `/tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task` directory.

Check whether there are too many pending compaction requests.

Step 7 Check whether the sum of the values of `rocksdb.mem-table-flush-pending` and `rocksdb.compaction-pending` in TaskManager monitoring logs (keyword `RocksDBMetricPrint`) is greater than `n` times the number of RocksDB threads (`metrics.reporter.alarm.job.alarm.rocksdb.background.jobs.multiplier`, 2 by default). If it is, you can increase the number of RocksDB threads.

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to **Step 8**.

Table 7-109 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.thread.num	<ul style="list-style-type: none"> - 2 - 4: enables SPINNING_DISK_OPTIMIZE and HIGH_MEMORY. 	<ul style="list-style-type: none"> - Number of flush threads. Increase the number of threads to quickly flush memory data to disks. - When the number of threads is increased, the number of vCores also needs to be increased. - 2 to 10 are recommended.

- If no, go to **Step 9**.

Step 8 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.368 ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the estimated pending compaction size of RocksDB for a job continuously exceeds the threshold. This alarm is cleared when the estimated pending compaction size of RocksDB for the job is less than or equal to the threshold.

The threshold of the estimated pending compaction size is the smaller value of the following two parameters:

- **state.backend.rocksdb.soft-pending-compaction-bytes-limit**. The default value is **64GB**.
- **state.backend.rocksdb.hard-pending-compaction-bytes-limit**. The default value is **256GB**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45647	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes


The estimated compaction data size of RocksDB is too large.

Handling Procedure

Check TaskManager logs for the estimated compaction data size and collect logs.

- Step 1** Log in to FusionInsight Manager as a user who has the FlinkServer management permission.
- Step 2** Choose **O&M > Alarm > Alarms > ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.
- Step 3** Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.
- Step 4** Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-230 Application ID of a job



- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-231 Clicking Logs

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to [Step 7](#).

Figure 7-232 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

Figure 7-233 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https/	0	Logs
container_0009_01_000001	https/	0	Logs

Showing 1 to 2 of 2 entries

NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the `/tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task` directory.

Check whether the estimated compaction data size of RocksDB is too large.

Step 7 Check whether the value of `rocksdb.estimate-pending-compaction-bytes` (unit: byte) in TaskManager monitoring logs (keyword `RocksDBMetricPrint`) is greater than or equal to the `soft/hard-pending-compaction` size (values of `state.backend.rocksdb.soft-pending-compaction-bytes-limit` and `state.backend.rocksdb.hard-pending-compaction-bytes-limit`).

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to [Step 8](#).

Table 7-110 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.soft-pending-compaction-bytes-limit	64GB	<ul style="list-style-type: none">- When the pending compaction size exceeds the threshold, the write traffic is limited.- 64GB to 512GB are recommended.
state.backend.rocksdb.hard-pending-compaction-bytes-limit	256GB	<ul style="list-style-type: none">- When the pending compaction size exceeds the threshold, write operations are stopped.- 64GB to 512GB are recommended.

- If no, go to [Step 9](#).

Step 8 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Step 9 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.369 ALM-45648 RocksDB Frequently Encounters Write-Stopped

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when RocksDB for a job continuously encounters the **is-write-stopped** state. This alarm is cleared when RocksDB for the job no longer or does not continuously encounter the **is-write-stopped** state within an alarm reporting interval.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45648	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes

The possible causes are as follows:

- There are too many MemTables and **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** is generated.
- There are too many SST files at level 0, and **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** is generated.
- The estimated compaction size exceeds the threshold, and **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** is generated.

Handling Procedure

Check whether there are too many MemTables.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 2 In the alarm list, check whether **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Handle the alarm by following the instructions provided in section **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45643 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check whether the number of SST files at level 0 is too large.

Step 5 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 6 In the alarm list, check whether **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

Step 7 Handle the alarm by following the instructions provided in section **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**.

Step 8 After ALM-45644 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Check whether the estimated compaction size exceeds the threshold.

Step 9 In the alarm list, check whether **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 10](#).
- If no, go to [Step 12](#).

Step 10 Handle the alarm by following the instructions provided in section **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 11 After ALM-45647 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

Collect fault information.

Step 12 Log in to Manager as a user who has the management permission for the current Flink job.

Step 13 Choose **O&M > Alarm > Alarms > ALM-45648 RocksDB Frequently Encounters Write-Stopped**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 14 Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 15 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-234 Application ID of a job

ID	User	QueueUser	Name
application			

- If yes, go to **Step 16**.
- If no, go to **Step 18**.

Step 16 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-235 Clicking Logs

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs.

Figure 7-236 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

Figure 7-237 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https/	0	Logs
container_0009_01_000001	https/	0	Logs

Showing 1 to 2 of 2 entries

 NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 17 View the job logs to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 18 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the */tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task* directory.

Step 19 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.370 ALM-45649 P95 Latency of RocksDB Get Requests Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the P95 latency of RocksDB Get requests exceeds the threshold (**metrics.reporter.alarm.job.alarm.rocksdb.get.micros.threshold**, 50000 microseconds by default). This alarm is cleared when the P95 latency of RocksDB Get requests is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45649	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes

The possible causes are as follows:

- There are too many SST files at level 0, causing slow queries. In addition, **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** is generated.
- The cache hit ratio is lower than 60%, causing frequent swap-ins and swap-outs of the block cache.

Handling Procedure

Check whether the number of SST files at level 0 is too large.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 2 In the alarm list, check whether **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Handle the alarm by following the instructions provided in section **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45644 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.

- If no, go to [Step 5](#).

Check the cache hit ratio in TaskManager logs and collect logs.

Step 5 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 6 Choose **O&M > Alarm > Alarms > ALM-45649 P95 Latency of RocksDB Get Requests Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 7 Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 8 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-238 Application ID of a job

ID	User	QueueUser	Name
application_...			

- If yes, go to [Step 9](#).
- If no, go to [Step 10](#).

Step 9 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-239 Clicking Logs


Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://...	Logs

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to [Step 11](#).

Figure 7-240 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://...	Logs

Figure 7-241 Clicking Logs



NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 10 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the `/tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task` directory.

Check whether the cache hit ratio is too low.

Step 11 Check the values of **rocksdb.block.cache.hit** (cache hit) and **rocksdb.block.cache.miss** (cache miss) in TaskManager monitoring logs (keyword **RocksDBMetricPrint**). Calculate the hit ratio using the following formula and check whether it is less than 60%:

$$\text{rocksdb.block.cache.hit} / (\text{rocksdb.block.cache.hit} + \text{rocksdb.block.cache.miss})$$

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to **Step 12**.

Table 7-111 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.block.cache-size	<ul style="list-style-type: none"> - 8MB - 256MB: enables SPINNING_DISK_OPTIMIZED_HIGH_MEM. 	<ul style="list-style-type: none"> - Cache size - 8MB to 1GB are recommended.
state.backend.rocksdb.block.blocksize	<ul style="list-style-type: none"> - 4KB - 128KB: enables SPINNING_DISK_OPTIMIZED_HIGH_MEM. 	<ul style="list-style-type: none"> - Block size - 4KB to 256KB are recommended.
state.backend.rocksdb.use-bloom-filter	false	<ul style="list-style-type: none"> - Whether to speed up indexing. If it is true, each new SST file will contain a Bloom filter. - true is recommended.

- If no, go to [Step 13](#).

Step 12 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 13](#).

Step 13 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.371 ALM-45650 P95 Latency of RocksDB Write Requests Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the P95 latency of RocksDB write requests exceeds the threshold (**metrics.reporter.alarm.job.alarm.rocksdb.write.micros.threshold**, 50000 microseconds by default). This alarm is cleared when the P95 latency of RocksDB write requests is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45650	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

The checkpoint performance of Flink jobs is affected. There is no impact on the FlinkServer.

Possible Causes

The possible causes are as follows:

- There are too many MemTables. As a result, write traffic is limited or write stops, and **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** is generated.
- There are too many SST files at level 0, and **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** is generated.
- The estimated compaction size exceeds the threshold, and **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** is generated.

Handling Procedure

Check whether write traffic limiting or write stop is caused due to too many MemTables.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 2 In the alarm list, check whether **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Handle the alarm by following the instructions provided in section **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45643 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check whether the number of SST files at level 0 is too large.

Step 5 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 6 In the alarm list, check whether **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

Step 7 Handle the alarm by following the instructions provided in section **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**.

Step 8 After ALM-45644 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Check whether the estimated compaction size exceeds the threshold.

Step 9 In the alarm list, check whether **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to [Step 10](#).
- If no, go to [Step 12](#).

Step 10 Handle the alarm by following the instructions provided in section **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 11 After ALM-45647 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 12](#).

Collect fault information.

Step 12 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 13 Choose **O&M > Alarm > Alarms > ALM-45650 P95 Latency of RocksDB Write Requests Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 14 Choose **Cluster > Services > Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 15 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 7-242 Application ID of a job

Show 20 entries			
ID	User	QueueUser	Name
application			

- If yes, go to **Step 16**.
- If no, go to **Step 18**.

Step 16 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 7-243 Clicking Logs

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs.

Figure 7-244 Clicking the ID in the Attempt ID column

Attempt ID	Started	Node	Logs	
appattempt_1_0009_000001	Mon Aug 30 14:51:33 +0800 2021	https://-	Logs	0

Showing 1 to 1 of 1 entries

Figure 7-245 Clicking Logs

Container ID	Node	Container Exit Status	Logs
container_0009_01_000002	https://-	0	Logs
container_0009_01_000001	https://-	0	Logs

Showing 1 to 2 of 2 entries

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission. Choose **Cluster > Services > Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 17 View the job logs to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 18 On Manager, choose **Cluster > Services > HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities > Browse the file system**, and download logs in the */tmp/logs/Username/bucket-logs-tfile/Last four digits of the task application ID/Application ID of the task* directory.

Step 19 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.372 ALM-45652 Flink Service Unavailable

This section applies to MRS 3.3.0 or later.

Alarm Description

The alarm module checks the Flink status every 60 seconds. This alarm is generated when the Flink service is unavailable. This alarm is cleared when the Flink service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45652	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the job for which the alarm is generated.

Impact on the System

Flink jobs cannot be submitted with FlinkServer and the Flink client.

Possible Causes

The ZooKeeper, HDFS, Yarn, KrbServer, or DBService service on which Flink depends is unavailable.

Handling Procedure

Check whether the ZooKeeper service on which Flink depends is abnormal.

Step 1 Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**.

Step 2 In the alarm list, check whether "ALM-13000 ZooKeeper Service Unavailable" exists.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

Step 3 Handle the alarm by referring to "ALM-13000 ZooKeeper Service Unavailable."

Step 4 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check whether the HDFS service on which Flink depends is abnormal.

Step 5 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 6 In the alarm list, check whether "ALM-14000 HDFS Service Unavailable" exists.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

Step 7 Handle the alarm by referring to "ALM-14000 HDFS Service Unavailable."

Step 8 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Check whether the Yarn service on which Flink depends is abnormal.

Step 9 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 10 In the alarm list, check whether "ALM-18000 Yarn Service Unavailable" exists.

- If yes, go to [Step 11](#).
- If no, go to [Step 13](#).

Step 11 Handle the alarm by referring to "ALM-18000 Yarn Service Unavailable."

Step 12 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 13](#).

Check whether the KrbServer service on which Flink depends is abnormal.

Step 13 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 14 In the alarm list, check whether "ALM-25500 KrbServer Service Unavailable" exists.

- If yes, go to [Step 15](#).
- If no, go to [Step 17](#).

Step 15 Handle the alarm by referring to "ALM-25500 KrbServer Service Unavailable."

Step 16 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 17](#).

Check whether the DBService service on which Flink depends is abnormal.

Step 17 On FusionInsight Manager, choose **O&M > Alarm > Alarms**.

Step 18 In the alarm list, check whether "ALM-27001 DBService Service Unavailable" exists.

- If yes, go to [Step 19](#).
- If no, go to [Step 21](#).

Step 19 Handle the alarm by referring to "ALM-27001 DBService Service Unavailable."


Step 20 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 21](#).

Collect fault information.

Step 21 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 22 Expand the **Service** drop-down list, and select **Flink** for the target cluster.

Step 23 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 24 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.373 ALM-45653 Invalid Flink HA Certificate File

This section applies to MRS 3.3.0 or later.

Alarm Description

Flink checks whether the HA certificate file is valid (whether the certificate exists and whether its format is correct) in the first health check or at 01:00:00 every day. This alarm is generated when the certificate file is invalid. This alarm is automatically cleared when the certificate file becomes valid again.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45653	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

FlinkServer in active/standby mode cannot provide services for external systems, and Flink jobs cannot be submitted on the FlinkServer.

Possible Causes

The HA certificate file is invalid.

Handling Procedure

View alarm information.


Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45653 Invalid Flink HA Certificate File**, view **Location**, obtain the name of the host for which the alarm is generated, and click the host name to view its IP address.

Check whether the HA certificate file in the system is valid.

Step 2 Log in to the host for which the alarm is generated as user **omm**.

- Step 3** Run the `cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/ha/local/cert` command to go to the directory where the HA certificate is stored.
- Step 4** Run the `ls -l` command to check whether the `server.crt` file exists.
- If yes, go to [Step 5](#).
 - If no, go to [Step 6](#).
- Step 5** Run the `openssl x509 -in server.crt -text -noout` command and check whether the command output is normal.
- If yes, go to [Step 9](#).
 - If no, go to [Step 6](#).
- Step 6** Run the `cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/flink/sbin` command to go to the Flink script directory.
- Step 7** Run the `sh proceed_ha_ssl_cert.sh` command to generate a new HA certificate. Then, check whether the alarm is cleared 1 minute later.
- If yes, go to [Step 8](#).
 - If no, go to [Step 9](#).
- Step 8** Check whether this alarm is generated again during periodic system check.
- If yes, go to [Step 9](#).
 - If no, no further action is required.

Collect fault information.

- Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.
- Step 10** Expand the **Service** drop-down list, and select **Flink** for the target cluster.
- Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 12** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.374 ALM-45654 Flink HA Certificate Is About to Expire

This section applies to MRS 3.3.0 or later.

Alarm Description

Flink checks whether the HA certificate file is about to expire in the first health check or at 01:00:00 every day. This alarm is generated when the remaining validity period is less than or equal to 30 days. This alarm is automatically cleared when the remaining validity period is greater than 30 days.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45654	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the certificate expires, the HA function of the FlinkServer in active/standby mode is affected. Flink jobs cannot be submitted on the FlinkServer. For FlinkServers in dual-active mode, the HA function is not affected.

Possible Causes

The HA certificate is about to expire.

Handling Procedure


View alarm information.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45654 Flink HA Certificate Is About to Expire**, view **Location**, obtain the name of the host for which the alarm is generated, and click the host name to view its IP address.

Check whether the HA certificate file in the system is valid. If it is not, generate a new one.

- Step 2** Log in to the host for which the alarm is generated as user **omm**.
- Step 3** Run the `cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/ha/local/cert` command to go to the directory where the HA certificate is stored.
- Step 4** Run the `openssl x509 -noout -text -in server.crt` command to query the effective time and due time of the HA certificate.
- Step 5** Perform [Step 6](#) to [Step 7](#) during off-peak hours to update the certificate file as needed.
- Step 6** Run the `cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/flink/sbin` command to go to the Flink script directory.
- Step 7** Run the `sh proceed_ha_ssl_cert.sh` command to generate a new HA certificate. Then, check whether the alarm is cleared 1 minute later.
- If yes, go to [Step 9](#).
 - If no, go to [Step 8](#).
- Step 8** On the node where the standby FlinkServer instance is located, repeat [Step 6](#) to [Step 7](#). Then, check whether the alarm is cleared 1 minute later.
- If yes, go to [Step 9](#).
 - If no, go to [Step 10](#).
- Step 9** Check whether this alarm is generated again during periodic system check.
- If yes, go to [Step 10](#).
 - If no, no further action is required.

Collect fault information.

- Step 10** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.
- Step 11** Expand the **Service** drop-down list, and select **Flink** for the target cluster.
- Step 12** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 13** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.375 ALM-45655 Flink HA Certificate File Has Expired

This section applies to MRS 3.3.0 or later.

Alarm Description

Flink checks whether the HA certificate file has expired in the first health check or at 01:00:00 every day. This alarm is generated when the HA certificate has expired. This alarm is automatically cleared when the certificate file becomes valid again.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45655	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

FlinkServer in active/standby mode cannot provide services for external systems, and Flink jobs cannot be submitted on the FlinkServer.

Possible Causes

The HA certificate file has expired.

Handling Procedure

View alarm information.


Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms > ALM-45655 Flink HA Certificate File Has Expired**, view **Location**, obtain the name of the host for which the alarm is generated, and click the host name to view its IP address.

Check whether the HA certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the host for which the alarm is generated as user **omm**.

- Step 3** Run the `cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/ha/local/cert` command to go to the directory where the HA certificate is stored.
- Step 4** Run the `openssl x509 -noout -text -in server.crt` command to query the effective time and due time of the HA certificate and check whether the HA certificate file is valid.
- If yes, go to [Step 9](#).
 - If no, go to [Step 5](#).
- Step 5** Run the `cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/flink/sbin` command to go to the Flink script directory.
- Step 6** Run the `sh proceed_ha_ssl_cert.sh` command to generate a new HA certificate. Then, check whether the alarm is cleared 1 minute later.
- If yes, go to [Step 8](#).
 - If no, go to [Step 7](#).
- Step 7** On the node where the standby FlinkServer instance is located, repeat [Step 5](#) to [Step 6](#). Then, check whether the alarm is cleared 1 minute later.
- If yes, go to [Step 8](#).
 - If no, go to [Step 9](#).
- Step 8** Check whether this alarm is generated again during periodic system check.
- If yes, go to [Step 9](#).
 - If no, no further action is required.

Collect fault information.

- Step 9** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.
- Step 10** Expand the **Service** drop-down list, and select **Flink** for the target cluster.
- Step 11** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 12** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.376 ALM-45736 Guardian Service Unavailable

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

The alarm module checks the Guardian service status every 60 seconds. This alarm is generated if Guardian is unavailable.

This alarm is cleared after Guardian recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45736	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Guardian cannot work properly, and OBS cannot be accessed.

Possible Causes

- The HDFS service on which the Guardian service depends is abnormal.
- The TokenServer role instance is abnormal.

Handling Procedure

Check the HDFS service status.

Step 1 Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. On the page that is displayed, check whether "ALM-14000 HDFS Service Unavailable" is reported.

- If yes, go to [Step 2](#).
- If no, go to [Step 3](#).

Step 2 Clear this alarm according to the alarm help.

After the alarm is cleared, wait a few minutes and check whether the alarm GuardianService Unavailable is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

Check all TokenServer instances.

Step 3 Log in to the node where the TokenServer instance resides as user **omm** and run the **ps -ef|grep "guardian.token.server.Server"** command to check whether the TokenServer process exists on the node.

- If yes, go to [Step 5](#).
- If no, restart the faulty TokenServer instance and go to [Step 4](#).


Step 4 In the alarm list, check whether the alarm "Guardian Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M > Log > Download**.

Step 6 Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.377 ALM-45737 TokenServer Heap Memory Usage Exceeds the Threshold

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

The system checks the heap memory usage of the TokenServer service every 60 seconds. This alarm is generated when the heap memory usage of the TokenServer instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times.

This alarm is automatically cleared when the system detects that the heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45737	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the heap memory of the Guardian TokenServer instance overflows, OBS cannot be accessed.

Possible Causes

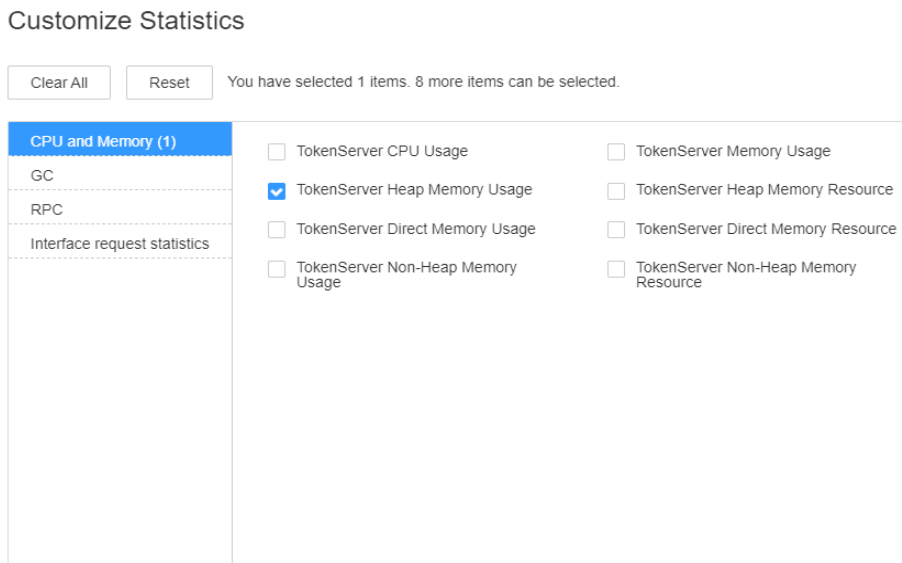
The heap memory of the TokenServer instance is overused or the heap memory is inappropriately allocated.

Handling Procedure

Check heap memory usage.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms > ALM-45737 TokenServer Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > TokenServer Heap Memory Usage**. Then click **OK**.

Figure 7-246 TokenServer Heap Memory Usage



Step 3 Check whether the heap memory used by TokenServer reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster > Services > Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, choose **TokenServer > Instance Configuration**. Click **All Configurations**, and choose **TokenServer > System**. Set **-Xmx** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for TokenServer cannot meet the heap memory required by the TokenServer process. You are advised to change the value of **-Xmx** in **GC_OPTS** to twice that of the heap memory used by TokenServer. You can change the value based on the actual service scenario. Refer to **Step 2** to view the TokenServer heap memory usage.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.


NOTICE

During service or instance restart, Guardian may fail to be accessed and jobs cannot access OBS.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M > Log > Download**.

Step 7 Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.378 ALM-45738 TokenServer Direct Memory Usage Exceeds the Threshold

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

The system checks the direct memory usage of the TokenServer service every 60 seconds. This alarm is generated when the direct memory usage of the TokenServer instance exceeds the threshold (80% of the maximum memory) for five consecutive times.

This alarm is automatically cleared when the system detects that the TokenServer direct memory usage is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45738	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the direct memory of the Guardian TokenServer instance overflows, OBS cannot be accessed.

Possible Causes

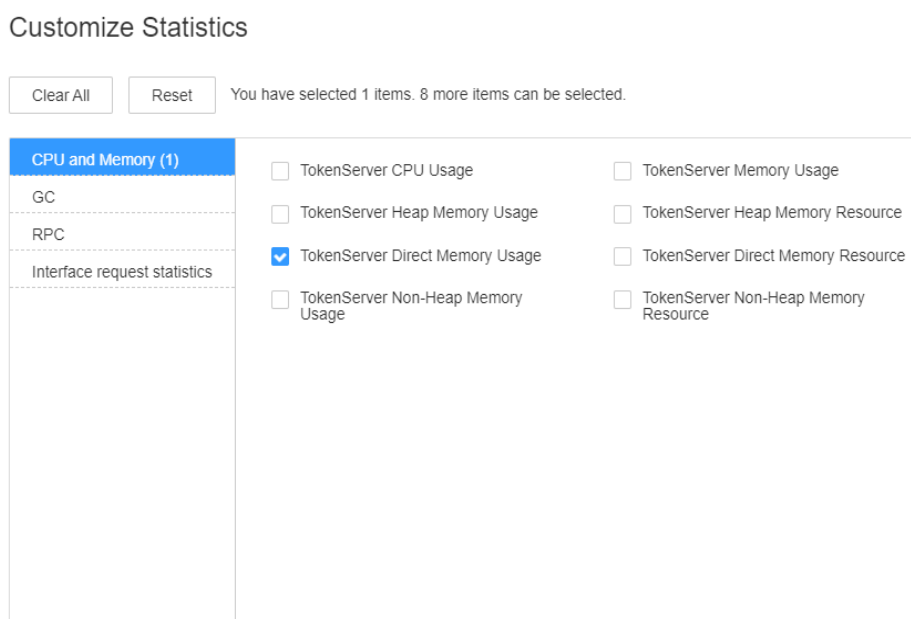
The direct memory of the TokenServer process is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms > ALM-45738 TokenServer Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > TokenServer Direct Memory Usage**. Then click **OK**.

Figure 7-247 TokenServer Direct Memory Usage



- Step 3** Check whether the direct memory used by TokenServer reaches the threshold (80% of the maximum direct memory by default).
- If yes, go to [Step 4](#).
 - If no, go to [Step 6](#).
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, choose **TokenServer > Instance Configuration**. Click **All Configurations**, and choose **TokenServer > System**. Set **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

 **NOTE**


If this alarm is generated, the direct memory configured for TokenServer cannot meet the direct memory required by the TokenServer process. You are advised to check the direct memory usage of TokenServer and change the value of **-XX:MaxDirectMemorySize** in **GC_OPTS** to the twice of the direct memory used by TokenServer. You can change the value based on the actual service scenario. Refer to [Step 2](#) to view the TokenServer direct memory usage.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 6](#).

NOTICE

During service or instance restart, Guardian may fail to be accessed and jobs cannot access OBS.

Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.379 ALM-45739 TokenServer Non-Heap Memory Usage Exceeds the Threshold

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

The system checks the heap memory usage of the TokenServer service every 60 seconds. This alarm is generated when the non-heap memory usage of the TokenServer instance exceeds the threshold (80% of the maximum memory) for five consecutive times.

This alarm is automatically cleared when the system detects that the non-heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45739	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the non-heap memory of the Guardian TokenServer instance overflows, OBS cannot be accessed.

Possible Causes

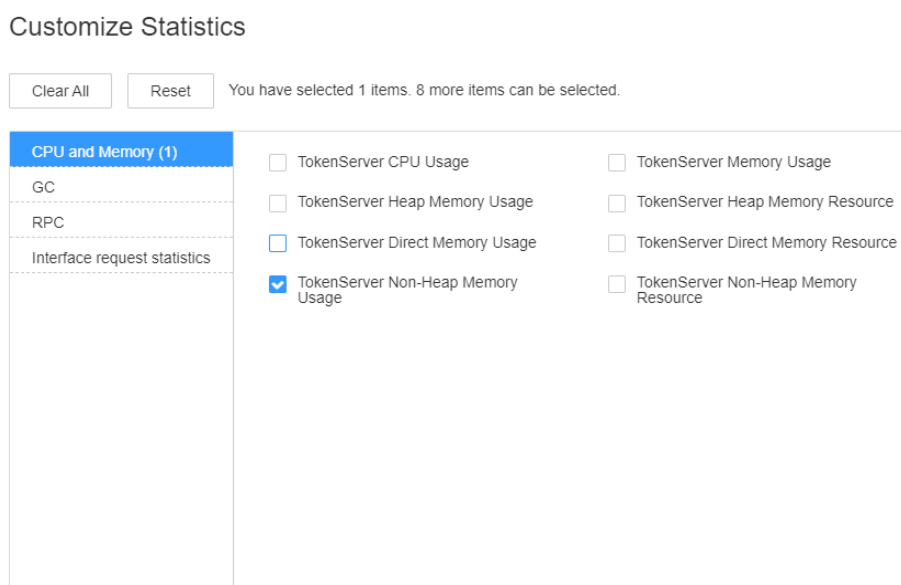
The non-heap memory of the TokenServer instance is overused or the non-heap memory is inappropriately allocated.

Handling Procedure

Check non-heap memory usage.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms > ALM-45739 TokenServer Non-Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize > CPU and Memory > TokenServer Non-Heap Memory Usage**. Then click **OK**.

Figure 7-248 TokenServer Non-Heap Memory Usage



- Step 3** Check whether the non-heap memory used by TokenServer reaches the threshold (80% of the maximum non-heap memory by default).
- If yes, go to **Step 4**.
 - If no, go to **Step 6**.
- Step 4** On FusionInsight Manager, choose **Cluster > Services > Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, choose **TokenServer > Instance Configuration**. Click **All Configurations**, and choose **TokenServer > System**. Set **-XX:MaxPermSize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE


If this alarm is generated, the non-heap memory size configured for the TokenServer instance cannot meet the non-heap memory required by the TokenServer process. You are advised to change the value of **-XX:MaxPermSize** in **GC_OPTS** to twice that of the current non-heap memory size or change the value based on site requirements.

- Step 5** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 6](#).

NOTICE

During service or instance restart, Guardian may fail to be accessed and jobs cannot access OBS.

Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M > Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.380 ALM-45740 TokenServer GC Duration Exceeds the Threshold

 **NOTE**

This section applies only to MRS 3.1.5 or later.

Alarm Description

The system checks the GC duration of the TokenServer process every 60 seconds. This alarm is generated when the GC duration of the TokenServer process exceeds the threshold (12 seconds by default) for five consecutive times.

This alarm is automatically cleared when the system detects that the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45740	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

TokenServer responds slowly, and OBS cannot be accessed.


Possible Causes

The heap memory of the TokenServer process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms > ALM-45740 TokenServer GC Duration Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.
- Step 2** On FusionInsight Manager, choose **Cluster > Services > Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize > GC > TokenServer GC Duration**. Then click **OK**.

Figure 7-249 TokenServer GC Duration

Step 3 Check whether the GC duration of the TokenServer process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to [Step 4](#).
- If no, go to [Step 6](#).

Step 4 On FusionInsight Manager, choose **Cluster > Services > Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, choose **TokenServer > Instance Configuration**. Click **All Configurations**, and choose **TokenServer > System**. Set **-Xmx** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for TokenServer cannot meet the heap memory required by the TokenServer process. You are advised to change the value of **-Xmx** in **GC_OPTS** to twice that of the heap memory used by TokenServer. You can change the value based on the actual service scenario. Refer to [Step 2](#) to view the TokenServer heap memory usage.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.


- If yes, no further action is required.
- If no, go to [Step 6](#).

NOTICE

During service or instance restart, Guardian may fail to be accessed and jobs cannot access OBS.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M > Log > Download**.

- Step 7** Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.381 ALM-45741 Failed to Call the ECS securitykey API

NOTE

This section applies only to MRS 3.2.1 or later.

Alarm Description

Guardian caches the temporary AK/SK of the ECS agency. When the cache does not exist or is about to expire, Guardian calls the securitykey API of ECS to update the AK/SK. This alarm is generated when calling to the API fails.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45741	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The task may fail to obtain the temporary AK/SK for accessing OBS. As a result, OBS cannot be accessed.

Possible Causes

- No ECS agency is bound to the cluster.
- An underlying interface of ECS is abnormal.

Handling Procedure

Check whether an agency is bound to the cluster.

Step 1 Log in to the MRS console.

Step 2 In the navigation pane on the left, choose **Clusters > Active Clusters**. On the page that is displayed, click the cluster name to go to its overview page. Then, check whether the cluster is bound to an agency in the O&M management area.

- If yes, go to [4](#).
- If no, go to [3](#).


Step 3 Click **Manage Agency**. On the page that is displayed, rebind the cluster to an agency. Then check whether the alarm is cleared a few minutes later.

- If yes, no further action is required.
- If no, go to [4](#).

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 5 Expand the **Service** drop-down list, and select **Guardian** for the target cluster.

Step 6 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.382 ALM-45742 Failed to Call the ECS Metadata API

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

When Guardian calls an IAM API to obtain the temporary AK/SK, it needs to first obtain related metadata via the ECS Metadata API. This alarm is generated when Guardian fails to call the Metadata API.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45742	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The task may fail to obtain the temporary AK/SK for accessing OBS. As a result, OBS cannot be accessed.


Possible Causes

An underlying interface of ECS is abnormal.

Handling Procedure

Collect fault information.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 2** Expand the **Service** drop-down list, and select **Guardian** for the target cluster.
- Step 3** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 4** Contact O&M personnel and provide the collected logs.
- End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.383 ALM-45743 Failed to Call the IAM API

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

This alarm is generated when Guardian fails to call the IAM API to obtain a temporary AK/SK.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45743	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System


The task may fail to obtain the temporary AK/SK for accessing OBS. As a result, OBS cannot be accessed.

Possible Causes

The IAM service is abnormal.

Handling Procedure

Collect fault information.

- Step 1** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 2** Expand the **Service** drop-down list, and select **Guardian** for the target cluster.
- Step 3** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 4** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

7.12.384 ALM-50201 Doris Service Unavailable

Alarm Description

The alarm module checks the Doris service status every 60 seconds. This alarm is generated when the alarm module detects that all FE and BE instances are abnormal.

This alarm is cleared when any FE or BE instance recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50201	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

FusionInsight Manager cannot be used to perform cluster operations on the Doris service, and Doris service functions are unavailable.

Possible Causes

The FE and BE instances are abnormal.

Handling Procedure

Restart the Doris service.

- Step 1** Log in to FusionInsight Manager and choose **Cluster > Services > Doris**.
- Step 2** On the page that is displayed, click **More** and select **Restart Service**. In the displayed dialog box, verify the password and click **OK** to restart the Doris service. After the service is started, go to [Step 3](#).

NOTICE

During the restart of the Doris service, the Doris service is unavailable and cannot provide services for external systems. Tasks connected to the Doris service fail.

- Step 3** On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the alarm list, check whether this alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 4](#).

Collect fault information.

- Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 5** Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.385 ALM-50202 FE CPU Usage Exceeds the Threshold

Alarm Description

The system checks the CPU usage of the FE instance every 30 seconds. The CPU usage has a default threshold. This alarm is generated when the CPU usage exceeds the threshold (**95%** by default) for multiple consecutive times (**3** by default).

This alarm is cleared when **Trigger Count** is **1** and the CPU usage is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the CPU usage is less than or equal to 85% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50202	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, select the name of the desired cluster, and choose **Doris > CPU and Memory > CPU Usage of FE (FE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

NOTE

Trigger Count specifies how many times the threshold can be hit before an alarm is generated.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.386 ALM-50203 FE Memory Usage Exceeds the Threshold

Alarm Description

The system checks the memory usage of the FE instance every 30 seconds. This alarm is generated when the memory usage exceeds the threshold (**95%** by default) for multiple consecutive times (**3** by default).

This alarm is cleared when **Trigger Count** is **1** and the memory usage is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the memory usage is less than or equal to 85% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50203	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

- The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, select the name of the desired cluster, and choose **Doris > CPU and Memory > Memory Usage of FE (FE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

 **NOTE**

Trigger Count specifies how many times the threshold can be hit before an alarm is generated.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.387 ALM-50205 BE CPU Usage Exceeds the Threshold

Alarm Description

The system checks the CPU usage of the BE instance every 30 seconds. This alarm is generated when the CPU usage exceeds the threshold (**95%** by default) for multiple consecutive times (**3** by default).

This alarm is cleared when **Trigger Count** is **1** and the CPU usage is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the CPU usage is less than or equal to 85% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50205	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > CPU and Memory > CPU Usage of BE (BE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

 **NOTE**

Trigger Count specifies how many times the threshold can be hit before an alarm is generated.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.388 ALM-50206 BE Memory Usage Exceeds the Threshold

Alarm Description

The system checks the memory usage of the BE instance every 30 seconds. This alarm is generated when the memory usage exceeds the threshold for multiple consecutive times (3 by default).

This alarm is cleared when **Trigger Count** is 1 and the memory usage is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the memory usage is less than or equal to 85% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50206	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > CPU and Memory > Memory Usage of BE (BE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

NOTE

Trigger Count specifies how many times the threshold can be hit before an alarm is generated.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.389 ALM-50207 Ratio of Connections to the FE MySQL Port to the Maximum Connections Allowed Exceeds the Threshold

Alarm Description

The system checks the number of MySQL port connections every 30 seconds. This alarm is generated when the ratio of the number of current connections to the maximum number of FE port connections exceeds the threshold (95% by default). The maximum number of FE port connections in the current cluster is specified by the **qe_max_connection** parameter. The default value is **1024**.

This alarm is cleared when the number of MySQL port connections on the FE node is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50207	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

- After the MySQL client is connected to Doris, the connection is not closed.
- A large number of services are concurrently connected to Doris.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M**, and click **Alarm > Thresholds** in the navigation pane on the left. Click the name of the desired cluster > **Doris > Connection > FE MySQL Port Connections (FE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

NOTE

If there are a large number of connections, ensure there are only necessary connections. Otherwise, the service performance may be degraded or even the service may be unavailable.

Step 4 Wait for 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.390 ALM-50208 Failures to Clear Historical Metadata Image Files Exceed the Threshold

Alarm Description

The system checks the number of failures to clear historical metadata image files on the FE node every 30 seconds. This alarm is generated when the number of failures exceeds the threshold (1 by default).

This alarm is cleared when the system detects that the number of failures is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50208	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Doris metadata occupies more and more disk space, which may cause service exceptions.

Possible Causes

The Doris service is abnormal.

Handling Procedure

Check whether the Doris service is normal.

Step 1 Log in to FusionInsight Manager and choose **Cluster > Services > Doris**.

Step 2 Check whether **Running Status** of the Doris service is **Normal**.

- If yes, go to [Step 4](#).
- If no, go to [Step 3](#).

Step 3 If the service process is not started, start it first and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Step 4 Check whether other Doris-related alarms are generated in the cluster. If yes, clear them by referring to the alarm help. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

You need to manually clear the alarm after the fault is rectified.

Related Information

None.

7.12.391 ALM-50209 Failures to Generate Metadata Image Files Exceed the Threshold

Alarm Description

The system checks the number of failures to generate metadata image files on the FE node every 30 seconds. This alarm is generated when the number of failures exceeds the threshold (1 by default).

This alarm is cleared when the system detects that the number of failures is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50209	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The non-master FE node cannot receive the latest metadata image file. As a result, the system reliability deteriorates.

Possible Causes

The Doris service is abnormal.

Handling Procedure

Check the Doris service status.

Step 1 Log in to FusionInsight Manager and choose **Cluster > Services > Doris**.

Step 2 Check whether **Running Status** of the Doris service is **Normal**.

- If yes, go to **Step 4**.
- If no, go to **Step 3**.

Step 3 If the service process is not started, start it first.

Step 4 Check whether other alarms are generated in the cluster. If yes, clear the alarms by referring to the alarm help. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

You need to manually clear the alarm after the fault is rectified.

Related Information

None.

7.12.392 ALM-50210 Maximum Compaction Score of All BE Nodes Exceeds the Threshold

Alarm Description

The system checks the maximum compaction score of all BE nodes every 30 seconds. This alarm is generated when the maximum compaction score exceeds the threshold (10 by default).

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50210	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Query or write may be delayed.

Possible Causes

The number of concurrent service requests is large in the cluster, or the compaction queue is small.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > Performance > Maximum compaction score of all BE nodes (BE)**.
- Step 2** Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.
- Step 3** Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.
- Step 4** Wait 2 minutes and check whether the alarm is cleared in the alarm list.
 - If yes, no further action is required.
 - If no, go to **Step 5**.
- Step 5** Choose **Cluster > Services > Doris > Configurations > All Configurations > BE(Role) > Customization**, add the **max_base_compaction_threads** parameter to **be.conf** with a value of **10**, and add the **max_cumu_compaction_threads** parameter with a value **20**.
- Step 6** Click **Save**. Click **Instances**, select the BE instances whose configuration has expired, click **More**, and select **Restart Instance** to restart the Doris BE instances.

NOTICE

During BE instance restart, the tasks running on BE nodes will fail. The tasks on BE nodes that are not restarted are not affected.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 9 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 10 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.393 ALM-50211 FE Queue Length of BE Periodic Report Tasks Exceeds the Threshold

Alarm Description

The system checks the queue length of each BE periodic report task on FE every 30 seconds. This alarm is generated when the queue length exceeds the threshold (10 by default). This value indicates the number of report tasks waiting on the master FE node. A large value indicates a poor FE processing capability.

This alarm is cleared when the system detects that the queue length of BE periodic report tasks on FE is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50211	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The processing capability of FE is insufficient, affecting the service query speed.

Possible Causes

The processing capability of the master FE node is insufficient due to a large number of concurrent service requests in the Doris cluster or insufficient memory for FE processes.

Handling Procedure

Check the GC duration.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the alarm list, view the role name and obtain the IP address of the instance in **Location** of the alarm whose ID is **50211**.

Step 2 Choose **Cluster > Services > Doris > Instances**, click the FE instance for which the alarm is generated, and click the **Chart** tab of the instance.

Select **JVM** from **Chart Category** on the left, and check whether **Accumulated GC duration of the old generation** of the FE process is greater than 3 seconds.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Choose **Cluster > Services > Doris > Configurations > All Configurations > FE(Role) > JVM**, and increase the value of **-Xmx** in **FE_GC_OPTS**. The default value is **8GB**.

- If this alarm is generated occasionally, increase the value by 0.5 times. If this alarm is generated frequently, double the parameter value.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 5 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > Queue > Queue Length of BE Periodic Report Tasks on the FE (FE)**.

Step 6 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 7 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 8 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 10 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.394 ALM-50212 Accumulated Old-Generation GC Duration of the FE Process Exceeds the Threshold

Alarm Description

The system checks the accumulated old-generation GC duration of the FE process every 30 seconds. This alarm is generated when the accumulated GC duration exceeds the threshold (3000 ms by default).

This alarm is cleared when the system detects that the accumulated old-generation GC duration of the FE process is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50212	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A long GC duration of the FE process may interrupt the services.

Possible Causes

The heap memory of the FE process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms**. In the alarm list, view the role name and obtain the IP address of the instance in **Location** of the alarm whose ID is **50212**.

Step 2 Choose **Cluster > Services > Doris > Instances**, click the FE instance for which the alarm is generated, and click the **Chart** tab of the instance.

Select **JVM** from **Chart Category** on the left, and check whether **Accumulated GC duration of the old generation** of the FE process is greater than 3 seconds.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

- Step 3** Choose **Cluster > Services > Doris > Configurations > All Configurations > FE(Role) > JVM**, and increase the value of **-Xmx** in **FE_GC_OPTS**. The default value is **8G**.
- If this alarm is occasionally generated, increase the value by 0.5 times. If this alarm is frequently generated, double the value.
 - In the case of large service volume and high service concurrency, you are advised to add instances.
- Step 4** Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 5](#).
- Collect fault information.**
- Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 6** Expand the **Service** drop-down list, and select **Doris** for the target cluster.
- Step 7** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 8** Contact O&M personnel and provide the collected logs.
- End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.395 ALM-50213 Number of Tasks Queuing in the FE Thread Pool for Interacting with BE Exceeds the Threshold

Alarm Description

The system checks the number of queuing tasks in the FE thread pool for interacting with BE every 30 seconds. This alarm is generated when the number of queuing tasks exceeds the threshold (10 by default). This FE thread pool is the working thread pool of ThriftServer. It is specified by **rpc_port** in the **fe.conf** file and is used to interact with BE.

This alarm is cleared when the system detects that the number of tasks queuing in the FE thread pool for interacting with BE is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50213	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The read and write of the Doris service slows down.

Possible Causes

There are a large number of concurrent service requests, causing too many queuing tasks.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > Queue > Number of tasks that are queuing in the thread pool for interaction between the FE and the BE (FE)**.
- Step 2** Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.
- Step 3** Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.
- Step 4** Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.396 ALM-50214 Number of Tasks Queuing in the FE Thread Pool for Task Processing Exceeds the Threshold

Alarm Description

The system checks the number of queuing tasks in the FE thread pool for processing tasks every 30 seconds. This alarm is generated when the number of queuing tasks exceeds the threshold (10 by default). This thread pool is used by the NIO MySQL Server to process tasks.

This alarm is cleared when the system detects that the number of tasks queuing in the FE thread pool for processing tasks is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50214	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The task execution of the entire system becomes slow and blocked.

Possible Causes

Large tasks may block the task execution of the queue.

Handling Procedure

Check the execution status of FE tasks.

Step 1 On FusionInsight Manager, choose **Cluster > Services > Doris**. Click the **Chart** tab, select **Connection** from **Chart Category** in the left pane, and view the **FE MySQL Port Connections** chart. If the number of connections is large, click **Instances**, select the FE instance, and click the **Chart** tab. Select **CPU and Memory** from **Chart Category** and view the **CPU Usage of FE** chart. If the CPU usage is high, check the **Time** field in FE audit log `/var/log/Bigdata/audit/doris/fe/fe.audit.log` to collect statistics on the average task duration. If the value is also high, the alarm is caused by large concurrent tasks.

Step 2 After connecting to Doris, run the following command to check whether the default value of **queryTimeout** is too large. The default value is **300** seconds.

show variables like 'query_timeout';

- If yes, go to [Step 3](#).
- If no, go to [Step 4](#).

Step 3 Run the following command to shorten the timeout period based on site requirements to block the tasks that take a long time:

set global query_timeout=xxx;

Step 4 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > Queue > Queue Length of Query Execution Thread Pool (BE)**.

Step 5 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 6 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 7 Wait 10 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 9 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 10 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.397 ALM-50215 Longest Duration of RPC Requests Received by Each FE Thrift Method Exceeds the Threshold

Alarm Description

The system checks the longest duration of RPC requests received by each FE Thrift method every 30 seconds. This alarm is generated when the longest duration exceeds the threshold (5000 ms by default).

This alarm is cleared when the longest duration of RPC requests received by each FE Thrift method is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50215	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A longer RPC duration indicates a higher performance load and slower network request processing, which may cause service congestion.

Possible Causes

- The network has a latency.
- There are too many concurrent large SQL tasks.

Handling Procedure

- Step 1** Log in to the host where the faulty node is deployed as user **root** and run **ping** *IP addresses of all Doris nodes* to check whether the peer host can be pinged.
- If yes, go to **Step 3**.
 - If no, go to **Step 2**.
- Step 2** Contact the network administrator to restore the network.
- Step 3** On FusionInsight Manager, choose **Cluster > Services > Doris**. Click the **Chart** tab, select **Connection** from **Chart Category** in the left pane, and view the **FE MySQL Port Connections** chart. If the number of connections is large, click **Instances**, select the FE instance, and click the **Chart** tab. Select **CPU and Memory** from **Chart Category** and view the **CPU Usage of FE** chart. If the CPU usage is high, check the **Time** field in FE audit log `/var/log/Bigdata/audit/doris/fe/fe.audit.log` to collect statistics on the average task duration. If the value is also high, the alarm is caused by large concurrent tasks.
- Step 4** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > Performance > Longest duration of RPC requests received by each method of the FE thrift interface (FE)**.
- Step 5** Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 6 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 7 Wait 10 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 10 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.398 ALM-50216 Memory Usage of the FE Node Exceeds the Threshold

Alarm Description

The system checks the memory usage of the FE node every 30 seconds. This alarm is generated when the memory usage exceeds the threshold (95% by default).

This alarm is cleared when the memory usage of the FE node falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50216	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Task execution and client connection to the FE are affected.

Possible Causes

The FE heap memory is too small.

Handling Procedure

Check the FE heap memory usage.

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, select the name of the desired cluster, and choose **Doris > CPU and Memory > Memory usage of the FE node (FE)**.
1. Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.
 2. Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.
- Step 2** Log in to the FE node for which the alarm is generated as user **omm**, run the **top** command to check the memory usage of processes, locate the process with high memory usage, and check whether the process belongs to the current service and is running properly.
- If yes, go to **Step 3**.
 - If no, isolate or stop the process, or adjust the memory size, and check whether the memory is released.
- Step 3** Restart the affected services or instances and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to **Step 4**.

NOTICE

- During the restart of the Doris service, the Doris service is unavailable and cannot provide services for external systems. Tasks connected to the Doris service fail.
- During instance restart, the tasks running on the nodes of the instance will fail. The tasks on instance nodes that are not restarted are not affected.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 5 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.399 ALM-50217 Heap Memory Usage of the FE Node Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the FE node every 30 seconds. This alarm is generated when the heap memory usage exceeds the threshold (95% by default).

This alarm is cleared when the heap memory usage of the FE node falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50217	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Task execution and client connection to the FE are affected.

Possible Causes

The FE heap memory is too small.

Handling Procedure

Check heap memory usage.

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, select the name of the desired cluster, and choose **Doris > CPU and Memory > Heap memory usage of the FE node (FE)**.

1. Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.
2. Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 2 On FusionInsight Manager, choose **Cluster > Services > Doris > FE > Configurations > All Configurations**, search for the **FE_GC_OPTS** parameter, increase the value of **-Xmx** as required, click **Save**, and click **OK**.

 NOTE

- If this alarm is generated, the heap memory configured for the current Doris instance is not enough for data transmission. You are advised to open the instance monitoring page, display the Doris heap memory resource status monitoring chart, and observe the change trend of the heap memory used by Doris in the monitoring chart. Then change the value of **-Xmx** to twice the current heap memory usage or to another value to meet site requirements.
- When setting the heap memory, you can set **-Xms** and **-Xmx** to approximately the same value to prevent performance deterioration caused by heap size adjustment after each GC.
- The sum of **-Xmx** and **XX:MaxPermSize** cannot be greater than the actual physical memory of the node server.

Step 3 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

NOTICE

- During the restart of the Doris service, the Doris service is unavailable and cannot provide services for external systems. Tasks connected to the Doris service fail.
 - During instance restart, the tasks running on the nodes of the instance will fail. The tasks on instance nodes that are not restarted are not affected.
-

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 5 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.400 ALM-50219 Length of the Queue in the Thread Pool for Query Execution Exceeds the Threshold

Alarm Description

The system checks the length of the waiting queue in the query execution thread pool every 30 seconds. This alarm is generated when the length exceeds the threshold (20 by default).

This alarm is cleared when the length of the waiting queue in the current query execution thread pool is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50219	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The task execution of the entire system becomes slow and blocked.

Possible Causes

Large tasks may block the task execution of the queue.

Handling Procedure

Check the execution status of tasks.

Step 1 On FusionInsight Manager, choose **Cluster > Services > Doris**. Click the **Chart** tab, select **Connection** from **Chart Category** in the left pane, and view the **FE MySQL**

Port Connections chart. If the number of connections is large, click **Instances**, select the FE instance, and click the **Chart** tab. Select **CPU and Memory** from **Chart Category** and view the **CPU Usage of FE** chart. If the CPU usage is high, check the **Time** field in FE audit log `/var/log/Bigdata/audit/doris/fe/fe.audit.log` to collect statistics on the average task duration. If the value is also high, the alarm is caused by large concurrent tasks.

Step 2 After connecting to Doris, run the following command to check the **queryTimeout** value of the system:

```
show variables like 'query_timeout';
```

If the value is too large, run the **set global query_timeout=xxx;** command to shorten the timeout interval and block tasks that last for a long time.

Step 3 Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > Queue > Queue Length of Query Execution Thread Pool (BE)**.

Step 4 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 5 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 6 Wait 10 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 8 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.401 ALM-50220 Error Rate of TCP Packet Receiving Exceeds the Threshold

Alarm Description

The system checks the rate of TCP packet receiving errors every 30 seconds. This alarm is generated when the error rate exceeds the threshold (5% by default).

This alarm is cleared when the error rate is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50220	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The task fails or data is lost.

Possible Causes

The network is faulty, so data cannot be sent.

Handling Procedure

Step 1 Log in to the host where the faulty node is deployed as user **root** and run **ping *IP addresses of all Doris nodes*** to check whether the peer host can be pinged.

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

Step 2 Contact the network administrator to restore the network.

Step 3 Wait for a while and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 5 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.402 ALM-50221 BE Data Disk Usage Exceeds the Threshold

Alarm Description

The system checks the usage of BE data disks every 30 seconds. This alarm is generated when the disk usage exceeds the threshold (95% by default).

This alarm is cleared when the system detects that the disk usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50221	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

New data fails to be written, and the task is interrupted.

Possible Causes

- The disk space of the cluster is full.
- Data skew occurs among BE nodes.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Expand the disk capacity of the node for which the alarm is generated.

Step 3 Go to [Step 4](#) if the expansion fails or the alarm persists after the expansion.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 5 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 6 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.403 ALM-50222 Disk Status of a Specified Data Directory on BE Is Abnormal

Alarm Description

The system checks the disk status of a specified data directory on BE every 30 seconds. This alarm is generated when the disk status is not **1** (**1** indicates the normal state and **0** indicates the abnormal state). This alarm is cleared when the disk status of the specified data directory on BE becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50222	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Service data may be unavailable, and data queries on the Doris client may fail.

Possible Causes

- The hard disk is faulty.
- The disk permissions are set incorrectly.

Handling Procedure

Check whether a disk alarm is generated.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms** and check whether **ALM-12014 Partition Lost** or **ALM-12033 Slow Disk Fault** exists.

- If yes, go to [Step 2](#).
- If no, go to [Step 4](#).

Step 2 Rectify the fault by referring to the handling procedure of **ALM-12014 Partition Lost** or **ALM-12033 Slow Disk Fault**. Then, check whether the alarm is cleared.

- If yes, go to [Step 3](#).
- If no, go to [Step 4](#).

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Modify disk permissions.

Step 4 Choose **O&M > Alarm > Alarms** and view **Location** and **Additional Information** of the alarm to obtain the location of the faulty disk.

Step 5 Log in to the node for which the alarm is generated as user **root**. Go to the directory where the faulty disk is located, and run the **ll** command to check whether the permission for the faulty disk is **711** and whether the user is **omm**.

- If yes, go to [Step 7](#).
- If no, go to [Step 6](#).

Step 6 Modify the permission of the faulty disk. For example, if the faulty disk is **data1**, run the following commands:

```
chown omm:wheel data1
```

```
chmod 711 data1
```

Step 7 Choose **Cluster > Services > Doris > Instances**, select this BE instance, click **More**, and select **Restart Instance**. Wait 5 minutes and check whether an alarm is generated.

- If no, no further action is required.
- If yes, go to [Step 8](#).

NOTICE

During BE instance restart, the tasks running on BE nodes will fail. The tasks on BE nodes that are not restarted are not affected.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

- Step 9** Expand the **Service** drop-down list, and select **Doris** and **OMS** for the target cluster.
- Step 10** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 11** Contact O&M personnel and provide the collected logs.
- End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.404 ALM-50223 Maximum Memory Required by BE Is Greater Than the Remaining Memory of the Machine

Alarm Description

The system checks whether the maximum memory required by BE is greater than the available memory every 30 seconds. This alarm is generated when the value is not **1** (**1** indicates that the maximum required memory is less than or equal to the available memory, and **0** indicates that the maximum required memory is greater than the available memory).

This alarm is cleared when the maximum required memory is less than or equal to the available memory.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50223	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A task may fail to apply for memory when running.

Possible Causes

Too much BE node memory has been occupied by other processes, or the maximum memory set for the BE service is too large.

Handling Procedure

Check whether the maximum memory set for the BE node is proper.

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > CPU and Memory > Relationship between the maximum memory size of the BE and the remaining memory size of the machine (BE)**.
- Step 2** Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.
- Step 3** Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.
- Step 4** Wait 2 minutes and check whether the alarm is automatically cleared.
 - If yes, no further action is required.
 - If no, go to [Step 5](#).
- Step 5** Log in to the BE node for which the alarm is generated as user **omm**, run the **top** command to check the memory usage of processes, locate the process with high memory usage, and check whether the process belongs to the current service and is running properly.
 - If yes, go to [Step 6](#).
 - If no, isolate or stop the process, or adjust the memory size, and check whether the memory is released.
- Step 6** Log in to the BE node for which the alarm is generated as user **omm** and run the **free -g** command to check the total memory and remaining memory in the system and estimate the memory usage.
- Step 7** On FusionInsight Manager, choose **Cluster > Services > Doris > Configurations > All Configurations > BE(Role) > Memory** and decrease the value of **mem_limit**. This parameter specifies the maximum memory allowed for BE. Then save the modification and restart the BE instance.

NOTICE

During BE instance restart, the tasks running on BE nodes will fail. The tasks on BE nodes that are not restarted are not affected.

Step 8 After the BE instance is restarted, wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 10 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.405 ALM-50224 Failures a Certain Task Type on BE Are Increasing

Alarm Description

The system checks whether the number of failed tasks of a certain type on BE is increasing every 30 seconds. This alarm is generated when the system detects that the value is not **1** (**1** indicates that the number of failed tasks of a certain type does not increase, and **0** indicates that the failed tasks of a certain type are increasing).

This alarm is cleared when the system detects that the number of failed tasks of a certain type on BE does not increase.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50224	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A task fails to be executed repeatedly in a certain scenario.

Possible Causes

A BE exception may occur. As a result, the number of failed tasks increases in a certain scenario.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

- Step 1** Log in to FusionInsight Manager, choose **O&M > Alarm > Thresholds**, click the name of the desired cluster, and choose **Doris > Exception > Check whether the number of failed tasks of a certain type increases (BE)**.
- Step 2** Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.
- Step 3** Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.
- Step 4** Wait 2 minutes and check whether the alarm is automatically cleared.
 - If yes, no further action is required.
 - If no, go to [Step 5](#).

Collect fault information.

- Step 5** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 6** Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

You need to manually clear the alarm after the fault is rectified.

Related Information

None.

7.12.406 ALM-50225 FE Instance Fault

Alarm Description

The system checks the FE process status every 30 seconds. This alarm is generated when the value is greater than **0** (**0** indicates that the FE process is normal and **1** indicates that the FE process is abnormal).

This alarm is cleared when the system detects that the FE process becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50225	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The FE instance is unavailable and cannot respond to client requests.

Possible Causes

The FE instance is faulty or restarted.

Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the alarm list, view the role name and obtain the IP address of the instance in **Location** of the alarm whose ID is **50225**.
 - Step 2** Choose **Cluster > Services > Doris > Instances**, click the FE instance for which the alarm is generated, and check whether **Running Status** of the instance is **Restoring**.
 - If yes, go to **Step 3**.
 - If no, go to **Step 4**.
 - Step 3** Wait 2 minutes and check whether the alarm is automatically cleared.
 - If yes, no further action is required.
 - If no, go to **Step 4**.
- Collect fault information.**
- Step 4** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
 - Step 5** Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.
 - Step 6** Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
 - Step 7** Contact O&M personnel and provide the collected logs.
- End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.407 ALM-50226 BE Instance Fault

Alarm Description

The system checks the BE process status every 30 seconds. This alarm is generated when the value is greater than **0** (**0** indicates that the BE process is normal and **1** indicates that the BE process is abnormal).

This alarm is cleared when the system detects that the BE process becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50226	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The BE instance is unavailable and cannot provide the data read and write functions.

Possible Causes

The BE instance is faulty or restarted.

Handling Procedure

- Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**. In the alarm list, view the role name and obtain the IP address of the instance in **Location** of the alarm whose ID is **50226**.
- Step 2** Choose **Cluster > Services > Doris > Instances**, click the BE instance for which the alarm is generated, and check whether **Running Status** of the instance is **Restoring**.
 - If yes, go to **Step 3**.
 - If no, go to **Step 4**.
- Step 3** Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 5 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.408 ALM-50401 Number of JobServer Jobs Waiting to Be Executed Exceeds the Threshold

Alarm Description

The system checks the number of jobs submitted to JobServer every 30 seconds. This alarm is generated when the number of jobs to be executed exceeds 800.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50401	Critical (default threshold: 900) Major (default threshold: 800)	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Too many JobServer jobs are detected in the queue. For details about the queue usage, see the **Additional Information** field of this alarm. The impacts are as follows:

1. When the number of JobServer jobs in the queue reaches the maximum (1000 by default), new jobs cannot be added.
2. Before the number of JobServer jobs in the queue reaches the maximum, new JobServer jobs cannot be submitted quickly. For example, it takes more time (even hours) to submit added jobs or add new jobs to Yarn.

Possible Causes

Too many jobs are submitted instantaneously.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **Cluster > Services > JobGateway**.

Step 2 Click the **Instances** tab, click **Add Instance**, and add JobServer instances based on the number of submitted jobs.

Step 3 After the instances are added, restart the JobGateway service.

NOTICE

The job functions of JobGateway will become unavailable during the service restart.

Step 4 Wait 5 minutes and check whether the alarm is automatically cleared.


If yes, no further action is required.

If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.

Step 6 Expand the **Service** drop-down list, and select **JobGateway** for the target cluster.

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.409 ALM-50402 JobGateway Service Unavailable

Alarm Description

The system checks the JobGateway service status every 60 seconds. This alarm is generated when the JobGateway service is abnormal.

This alarm is cleared when the JobGateway service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50402	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

No job submission operation can be performed on JobGateway in the cluster. The components that depend on JobGateway in the cluster will become faulty.


Possible Causes

The node where the JobGateway service locates is faulty.

Handling Procedure

- Step 1** Log in to FusionInsight Manager, choose **Cluster > Services > JobGateway**, and click the **Instance** tab. Check for JobServer or JobBalancer instances that are faulty or not started and view the host names of these instances.
- Step 2** On the **Alarm** page of FusionInsight Manager, check whether the **NodeAgent Process Is Abnormal** alarm is generated.
 - If yes, go to **Step 3**.
 - If no, go to **Step 6**.
- Step 3** Check whether the host name in the alarm information is the same as the host name in **Step 1**.
 - If yes, go to **Step 4**.
 - If no, go to **Step 6**.
- Step 4** Clear the alarm by following the instructions provided in **ALM-12006 NodeAgent Process Is Abnormal**.
- Step 5** In the alarm list, check whether alarm **JobGateway Service Unavailable** is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 6**.

Collect fault information.

- Step 6** On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
- Step 7** Expand the **Service** drop-down list, and select **JobGateway** for the target cluster.
- Step 8** Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
- Step 9** Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

7.12.410 ALM-12001 Audit Log Dump Failure (For MRS 2.x or Earlier)

Description

Cluster audit logs need to be dumped on a third-party server due to the local historical data backup policy. Audit logs can be successfully dumped if the dump server meets the configuration conditions. This alarm is generated when the audit log dump fails because the disk space of the dump directory on the third-party server is insufficient or a user changes the username, password, or dump directory of the dump server.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12001	Minor	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The system can only store a maximum of 50 dump files locally. If the fault persists on the dump server, the local audit log may be lost.

Possible Causes

- The network connection is abnormal.
- The username, password, or dump directory of the dump server does not meet the configuration conditions.
- The disk space of the dump directory is insufficient.

Procedure

Step 1 Check whether the username, password, and dump directory are correct.

1. Check on the dump configuration page of MRS Manager to see if they are correct.

- If yes, go to [Step 3](#).
 - If no, go to [Step 1.2](#).
2. Change the username, password, or dump directory, and click **OK**.
 3. Wait 2 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Reset the dump rule.

1. On MRS Manager, choose **System > Dump Audit Log**.
2. Reset dump rules, set the parameters properly, and click **OK**.
3. Wait 2 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 3](#).

Step 3 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

N/A

7.12.411 ALM-12002 HA Resource Abnormal (For MRS 2.x or Earlier)

Description

The high availability (HA) software periodically checks the Webservice floating IP addresses and databases of Manager. This alarm is generated when the HA software detects that the Webservice floating IP addresses or databases are abnormal.

This alarm is cleared when the HA software detects that the floating IP addresses or databases are normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12002	Major	Yes

Parameter

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
RESName	Specifies the resource for which the alarm is generated.

Impact on the System

If the Webservice floating IP addresses of Manager are abnormal, users cannot log in to or use Manager. If databases of Manager are abnormal, all core services and related service processes, such as alarms and monitoring functions, are affected.

Possible Causes

- The floating IP address is abnormal.
- The database is abnormal.

Procedure

Step 1 Check the floating IP address status of the active management node.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host address and resource name of the alarm.
2. Log in to the active management node. Run the following commands to switch the user:
sudo su - root
su - omm
3. Go to the ``${BIGDATA_HOME}/om-0.0.1/sbin/`` directory, run the **status-oms.sh** script to check whether the floating IP address of the active Manager is normal. View the command output, locate the row where **ResName** is **floatip**, and check whether the following information is displayed.
Example:

```
10-10-10-160 floatip Normal Normal Single_active
```

 - If yes, go to [Step 2](#).
 - If no, go to [Step 1.4](#).
4. Contact the O&M personnel to check whether the floating IP NIC exists.
 - If yes, go to [Step 2](#).

- If no, go to [Step 1.5](#).
5. Contact O&M personnel to rectify the NIC fault.
Wait 5 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Check the database status of the active and standby management nodes.

1. Log in to the active and standby management nodes, run the **sudo su - root** and **su - ommdba** commands to switch to user **ommdba**, and run the **gs_ctl query** command to check whether the following information is displayed in the command output.

Command output of the active management node:

```
Ha state:
LOCAL_ROLE: Primary
STATIC_CONNECTIONS: 1
DB_STATE: Normal
DETAIL_INFORMATION: user/password invalid
Senders info:
No information
Receiver info:
No information
```

Command output of the standby management node:

```
Ha state:
LOCAL_ROLE: Standby
STATIC_CONNECTIONS: 1
DB_STATE : Normal
DETAIL_INFORMATION: user/password invalid
Senders info:
No information
Receiver info:
No information
```

- If yes, go to [Step 2.3](#).
 - If no, go to [Step 2.2](#).
2. Contact the O&M personnel to check whether a network fault occurs and rectify the fault.
 - If yes, go to [Step 2.3](#).
 - If no, go to [Step 3](#).
 3. Wait 5 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 3](#).

Step 3 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.412 ALM-12004 OLdap Resource Abnormal (For MRS 2.x or Earlier)

Description

This alarm is generated when the Ldap resource in Manager is abnormal.

This alarm is cleared when the Ldap resource in Manager recovers and the alarm handling is complete.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12004	Major	Yes

Parameter

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The Manager authentication services are unavailable and cannot provide security authentication and user management functions for web upper-layer services. Users may be unable to log in to Manager.

Possible Causes

The LdapServer process in Manager is abnormal.

Procedure

Step 1 Check whether the LdapServer process in Manager is normal.

1. Log in to the active management node.
2. Run **ps -ef | grep slapd** to check whether the LdapServer resource process in the **`\${BIGDATA_HOME}/om-0.0.1/`** directory of the configuration file is running properly.

You can determine that the resource is normal as follows:

- a. Run `sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh` and find that **ResHAStatus** of the OLdap process is **Normal**.
- b. Run `ps -ef | grep slapd` and find that the slapd process occupies port 21750.
 - If yes, go to [Step 2](#).
 - If no, go to [Step 3](#).

Step 2 Run `kill -2 PID of the LdapServer process` and wait 20 seconds. The HA starts the OLdap process automatically. Check whether the status of the OLdap resource is normal.

- If yes, no further action is required.
- If no, go to [Step 3](#).

Step 3 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.413 ALM-12005 OKerberos Resource Abnormal (For MRS 2.x or Earlier)

Description

The alarm module monitors the status of the Kerberos resource in Manager. This alarm is generated when the Kerberos resource is abnormal.

This alarm is cleared when the alarm handling is complete and the Kerberos resource status recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12005	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The authentication services are unavailable and cannot provide security authentication functions for web upper-layer services. Users may be unable to log in to MRS Manager.

Possible Causes

The OLdap resource on which OKerberos depends is abnormal.

Procedure

Step 1 Check whether the OLdap resource on which OKerberos depends is abnormal in Manager.

1. Log in to the active management node.
2. Run the following command to check whether the OLdap resource managed by HA is normal:

```
sh ${BIGDATA_HOME}/OMSV100R001C00x8664/workspace0/ha/module/hacom/script/status_ha.sh
```

The OLdap resource is normal when the OLdap resource is in the **Active_normal** state on the active node and in the **Standby_normal** state on the standby node.

- If yes, go to [Step 3](#).
- If no, go to [Step 2](#).

Step 2 Resolve the problem by following the instructions in [ALM-12004 OLdap Resource Abnormal \(For MRS 2.x or Earlier\)](#). After the OLdap resource status recovers, check whether the OKerberos resource is normal.

- If yes, no further action is required.
- If no, go to [Step 3](#).

Step 3 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.414 ALM-12006 Node Fault (For MRS 2.x or Earlier)

Description

Controller checks the NodeAgent status every 30 seconds. This alarm is generated when Controller fails to receive the status report of a NodeAgent for three consecutive times.

This alarm is cleared when Controller can properly receive the status report of the NodeAgent.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12006	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Services on the node are unavailable.

Possible Causes

The network is disconnected, or the hardware is faulty.

Procedure

Step 1 Check whether the network is disconnected or the hardware is faulty.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host address of the alarm.
2. Log in to the active management node.
3. Run the following command to check whether the faulty node is reachable:
ping *IP address of the faulty host*

- a. If yes, go to [Step 2](#).
- b. If no, go to [Step 1.4](#).
4. Contact the O&M personnel to check whether the network is faulty.
 - If yes, go to [Step 2](#).
 - If no, go to [Step 1.6](#).
5. Rectify the network fault and check whether the alarm is cleared from the alarm list.
 - If yes, no further action is required.
 - If no, go to [Step 1.6](#).
6. Contact the O&M personnel to check whether a hardware fault (for example, a CPU or memory fault) occurs on the node.
 - If yes, go to [Step 1.7](#).
 - If no, go to [Step 2](#).
7. Repair the faulty components and restart the node. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.415 ALM-12007 Process Fault (For MRS 2.x or Earlier)

Description

The process health check module checks the process status every 5 seconds. This alarm is generated when the process health check module detects that the process connection status is Bad for three consecutive times.

This alarm is cleared when the process can be connected.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12007	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The service provided by the process is unavailable.

Possible Causes

- The instance process is abnormal.
- The drive space is insufficient.

Procedure

Step 1 Check whether the instance process is abnormal.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host name and service name of the alarm.
2. On the **Alarms** page, check whether the alarm **ALM-12006 Node Fault (For MRS 2.x or Earlier)** is generated.
If yes, go to **Step 1.3**.
If no, go to **Step 1.4**.
3. Handle the alarm by following the instructions in **ALM-12006 Node Fault (For MRS 2.x or Earlier)**.
4. Check whether the installation directory user, user group, and permission of the alarm role are correct. The correct user, user group, and the permission are **omm**, **ficommon**, and **750**, respectively.
 - If yes, go to **Step 1.6**.
 - If no, go to **Step 1.5**.
5. Run the following commands to set the permission to **750** and **User:Group** to **omm:ficommon**:
chmod 750 <folder_name>
chown omm:ficommon <folder_name>
6. Wait 5 minutes and check whether the ALM-12007 Process Fault alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2.1**.

Step 2 Check whether the disk space is insufficient.

1. On the MRS cluster details page, click the alarm management tab and check whether ALM-12017 Insufficient Disk Capacity is generated in the alarm list.
 - If yes, go to [Step 2.2](#).
 - If no, go to [Step 3](#).
2. Handle the alarm by following the instructions in [ALM-12017 Insufficient Disk Capacity \(For MRS 2.x or Earlier\)](#).
3. Wait 5 minutes and check whether the ALM-12017 Insufficient Disk Capacity alarm is cleared.
 - If yes, go to [Step 2.4](#).
 - If no, go to [Step 3](#).
4. Wait 5 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 3](#).

Step 3 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.416 ALM-12010 Manager Heartbeat Interruption Between the Active and Standby Nodes (For MRS 2.x or Earlier)

Description

This alarm is generated when the active Manager does not receive any heartbeat signal from the standby Manager within 7 seconds.

This alarm is cleared when the active Manager receives heartbeat signals from the standby Manager.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12010	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local Manager HA Name	Specifies a local Manager HA.
Peer Manager HA Name	Specifies a peer Manager HA.

Impact on the System

When the active Manager process is abnormal, an active/standby failover cannot be performed, and services are affected.

Possible Causes

The link between the active and standby Manager servers is abnormal.

Procedure

Step 1 Check whether the network between the active and standby Manager servers is normal.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the address of the standby Manager server.
2. Log in to the active management node.
3. Run the following command to check whether the standby Manager is reachable:
ping *heartbeat IP address of the standby Manager*
 - If yes, go to [Step 2](#).
 - If no, go to [Step 1.4](#).
4. Contact the O&M personnel to check whether the network is faulty.
 - If yes, go to [Step 1.5](#).
 - If no, go to [Step 2](#).
5. Rectify the network fault and check whether the alarm is cleared from the alarm list.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Log in to all master nodes in the cluster and run the following commands to find all **sedxxx** files and delete them:

```
find /srv/BigData/ -name "sed*"
```

```
find /opt -name "sed*"
```

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.417 ALM-12011 Data Synchronization Exception Between the Active and Standby Manager Nodes (For MRS 2.x or Earlier)

Description

This alarm is generated when the standby Manager fails to synchronize files with the active Manager.

This alarm is cleared when the standby Manager synchronizes files with the active Manager.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12011	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local Manager HA Name	Specifies a local Manager HA.
Peer Manager HA Name	Specifies a peer Manager HA.

Impact on the System

Because the configuration files on the standby Manager are not updated, some configurations will be lost after an active/standby switchover. Manager and some components may not run properly.

Possible Causes

The link between the active and standby Manager nodes is interrupted.

Procedure

- Step 1** Check whether the network between the active and standby Manager servers is normal.
1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the address of the standby Manager server.
 2. Log in to the active management node. Run the following command to check whether the standby Manager is reachable:
ping *IP address of the standby Manager*
 - If yes, go to **Step 2**.
 - If no, go to **Step 1.3**.
 3. Contact the O&M personnel to check whether the network is faulty.
 - If yes, go to **Step 1.4**.
 - If no, go to **Step 2**.
 4. Rectify the network fault and check whether the alarm is cleared from the alarm list.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

- Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.418 ALM-12012 NTP Service Abnormal (For MRS 2.x or Earlier)

Description

This alarm is generated when the NTP service on the current node fails to synchronize time with the NTP service on the active OMS node.

This alarm is cleared when the NTP service on the current node synchronizes time properly with the NTP service on the active OMS node.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12012	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The time on the node is inconsistent with that on other nodes in the cluster. Therefore, some MRS applications on the node may not run properly.

Possible Causes

- The NTP service on the current node cannot start properly.
- The current node fails to synchronize time with the NTP service on the active OMS node.
- The key value authenticated by the NTP service on the current node is inconsistent with that on the active OMS node.
- The time offset between the node and the NTP service on the active OMS node is large.

Procedure

Step 1 Check the NTP service on the current node.

1. Check whether the ntpd process is running on the node using the following method. Log in to the node for which the alarm is generated and run the **sudo su - root** command to switch to user **root**. Then run the following command to check whether the command output contains the ntpd process:

```
ps -ef | grep ntpd | grep -v grep
```

- If yes, go to [Step 2.1](#).
- If no, go to [Step 1.2](#).

2. Run **service ntp start** to start the NTP service.
3. Wait 10 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2.1](#).

Step 2 Check whether the current node can synchronize time properly with the NTP service on the active OMS node.

1. Check whether the node can synchronize time with the NTP service on the active OMS node based on additional information of the alarm.

If yes, go to [Step 2.2](#).

If no, go to [Step 3](#).

2. Check whether the synchronization with the NTP service on the active OMS node is faulty.

Log in to the node for which the alarm is generated, run the **sudo su - root** command to switch to user **root**, and run the **ntpq -np** command.

If an asterisk (*) exists before the IP address of the NTP service on the active OMS node in the command output, the synchronization is in normal state. The command output is as follows:

```
remote refid st t when poll reach delay offset jitter
```

```
=====
```

```
*10.10.10.162 .LOCL. 1 u 1 16 377 0.270 -1.562 0.014
```

If there is no asterisk (*) before the IP address of the NTP service on the active OMS node, as shown in the following command output, and the value of **refid** is **.INIT.**, the synchronization is abnormal.

```
remote refid st t when poll reach delay offset jitter
```

```
=====
```

```
10.10.10.162 .INIT. 1 u 1 16 377 0.270 -1.562 0.014
```

- If yes, go to [Step 2.3](#).
 - If no, go to [Step 3](#).
3. Rectify the fault, wait 10 minutes, and then check whether the alarm is cleared.

An NTP synchronization failure is usually related to the system firewall. If the firewall can be disabled, disable it and then check whether the fault is rectified. If the firewall cannot be disabled, check the firewall configuration policies and ensure that port **UDP 123** is enabled (you need to follow specific firewall configuration policies of each system).

- If yes, no further action is required.
- If no, go to [Step 3](#).

Step 3 Check whether the key value authenticated by the NTP service on the current node is consistent with that on the active OMS node.

Run **cat /etc/ntp.keys** to check whether the authentication code whose key value index is 1 is the same as the value of the NTP service on the active OMS node.

- If yes, go to [Step 4.1](#).
- If no, go to [Step 5](#).

Step 4 Check whether the time offset between the node and the NTP service on the active OMS node is large.

1. Check whether the time offset is large in additional information of the alarm.
 - If yes, go to [Step 4.2](#).
 - If no, go to [Step 5](#).

2. On the **Hosts** page, select the host of the node, and choose **More > Stop All Roles** to stop all the services on the node.

If the time on the alarm node is later than that on the NTP service of the active OMS node, adjust the time of the alarm node. After adjusting the time, choose **More > Start All Roles** to start the services on the node.

If the time on the alarm node is earlier than that on the NTP service of the active OMS node, wait until the time offset is due and adjust the time of the alarm node. After adjusting the time, choose **More > Start All Roles** to start the services on the node.

 **NOTE**

If you do not wait, data loss may occur.

3. Wait 10 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 5](#).

Step 5 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.419 ALM-12014 Device Partition Lost (For MRS 2.x or Earlier)

Description

This alarm is generated when the system detects that a partition to which service directories are mounted is lost (because the device is removed or goes offline, or the partition is deleted). The system checks the partition status periodically.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12014	Major	<ul style="list-style-type: none">• Yes: MRS 1.9.3.10 and later patch versions• No: MRS 2.x and earlier versions

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DirName	Specifies the directory for which the alarm is generated.
PartitionName	Specifies the device partition for which the alarm is generated.

Impact on the System

Service data fails to be written into the partition, and the service system runs abnormally.

Possible Causes

- The disk is removed.
- The disk is offline, or a bad sector exists on the disk.

Procedure

- Step 1** Go to the MRS cluster details page and choose **Alarms**.
- Step 2** In the real-time alarm list, click the row that contains the alarm.
- Step 3** In the **Alarm Details** area, obtain the values of **HostName**, **PartitionName**, and **DirName** from **Location**.
- Step 4** Check whether the disk corresponding to **PartitionName** on **HostName** is inserted to the correct server slot.
 - If yes, go to [Step 5](#).
 - If no, go to [Step 6](#).
- Step 5** Contact hardware engineers to remove the faulty disk.
- Step 6** Use PuTTY to log in to the **HostName** node where an alarm is reported and check whether there is a line containing **DirName** in the **/etc/fstab** file.
 - If yes, go to [Step 7](#).
 - If no, go to [Step 8](#).
- Step 7** Run the **vi /etc/fstab** command to edit the file and delete the line containing **DirName**.

- Step 8** Contact hardware engineers to insert a new disk. For details, see the hardware product document of the relevant model. If the faulty disk is in a RAID group, configure the RAID group. For details, see the configuration methods of the relevant RAID controller card.
- Step 9** Wait 20 to 30 minutes (The disk size determines the waiting time), and run the **mount** command to check whether the disk has been mounted to the **DirName** directory.
- If yes, perform **Step 10** for MRS 1.9.3.10 or later. For other versions, clear the alarm. No further action is required.
 - If no, perform **Step 11**.
- Step 10** Wait 2 minutes and check whether the alarm is automatically cleared.
- If yes, no further action is required.
 - If no, perform **Step 11**.
- Step 11** Collect fault information.
1. On MRS Manager, choose **System > Export Log**.
 2. Contact the O&M engineers and send the collected logs.
- End

Alarm Clearing

MRS 1.9.3.10 and later patch versions: After the fault is rectified, the system automatically clears the alarm.

MRS 2.x and earlier versions: After the fault is rectified, the system does not automatically clear the alarm. You need to clear the alarm.

Reference

None

7.12.420 ALM-12015 Device Partition File System Read-Only (For MRS 2.x or Earlier)

Description

This alarm is generated when the system detects that a partition to which service directories are mounted enters the read-only mode (due to a bad sector or a faulty file system). The system checks the partition status periodically.

This alarm is cleared when the system detects that the partition to which service directories are mounted exits from the read-only mode (because the file system is restored to read/write mode, the device is removed, or the device is formatted).

Attribute

Alarm ID	Alarm Severity	Auto Clear
12015	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DirName	Specifies the directory for which the alarm is generated.
PartitionName	Specifies the device partition for which the alarm is generated.

Impact on the System

Service data fails to be written into the partition, and the service system runs abnormally.

Possible Causes

The disk is faulty, for example, a bad sector exists.

Procedure

- Step 1** Go to the MRS cluster details page and choose **Alarms**.
- Step 2** In the real-time alarm list, click the row that contains the alarm.
- Step 3** In the **Alarm Details** area, obtain **HostName** and **PartitionName** from **Location**. **HostName** indicates the node for which the alarm is generated, and **PartitionName** indicates the partition of the faulty disk.
- Step 4** Contact hardware engineers to check whether the disk is faulty. If the disk is faulty, remove it from the server.
- Step 5** After the disk is removed, the system reports ALM-12014 Partition Lost. Handle the alarm by following the instructions in [ALM-12014 Device Partition Lost \(For MRS 2.x or Earlier\)](#). After the handling, the alarm is automatically cleared.

----End

Reference

None

7.12.421 ALM-12016 CPU Usage Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the CPU usage every 30 seconds and compares the check result with the default threshold. The CPU usage has a default threshold. This alarm is generated when the CPU usage exceeds the threshold for several times (configurable, 10 times by default) consecutively.

This alarm is cleared when the average CPU usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12016	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

- The alarm threshold or alarm hit number is improperly configured.
- The CPU configuration cannot meet service requirements. The CPU usage reaches the upper limit.

Procedure

Step 1 Check whether the alarm threshold or alarm hit number is properly configured.

1. Log in to MRS Manager and change the alarm threshold and alarm hit number based on CPU usage.
2. Choose **System > Threshold Configuration > Device > Host > CPU > CPU Usage > CPU Usage** and change the alarm threshold based on the actual CPU usage.
3. Choose **System > Threshold Configuration > Device > Host > CPU > CPU Usage > CPU Usage** and change **hit number** based on the actual CPU usage.

 **NOTE**

This option defines the alarm check phase. **Interval** indicates the alarm check period and **hit number** indicates the number of times when the CPU usage exceeds the threshold. An alarm is generated when the CPU usage exceeds the threshold for several times consecutively.

4. Wait 2 minutes and check whether the alarm is automatically cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Expand the system.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the address of the node.
2. Log in to the node for which the alarm is generated.
3. Run `cat /proc/stat | awk 'NR==1|awk '{for(i=2;i<=NF;i++)j+=Si;print "" 100 - ($5+$6) * 100 / j;}'` to check the system CPU usage.
4. If the CPU usage exceeds the threshold, expand the CPU capacity.
5. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 3](#).

Step 3 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.422 ALM-12017 Insufficient Disk Capacity (For MRS 2.x or Earlier)

Description

The system checks the host disk usage every 30 seconds and compares the actual disk usage with the threshold. The disk usage has a default threshold. This alarm is generated if the disk usage exceeds the threshold.

To change the threshold, choose **System > Threshold Configuration**.

This alarm is cleared when the host disk usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12017	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Service processes become unavailable.

Possible Causes

The disk configuration cannot meet service requirements. The disk usage reaches the upper limit.

Procedure

Step 1 Log in to MRS Manager and check whether the threshold is appropriate.

1. The default threshold is 90%. You can change the threshold to meet service requirements.
 - If yes, go to [Step 2](#).
 - If no, go to [Step 1.2](#).
2. Choose **System > Threshold Configuration** and change the alarm threshold based on the actual disk usage.
3. Wait 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 2](#).

Step 2 Check whether the disk is a system disk.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host name and disk partition information.
2. Log in to the node for which the alarm is generated.
3. Run the **df -h** command to check the system disk partition usage. Check whether the disk is mounted to any of the following directories by using the disk partition name obtained in [Step 2.1](#): **/**, **/boot**, **/home**, **/opt**, **/tmp**, **/var**, **/var/log**, **/boot**, and **/srv/BigData**.
 - If yes, the disk is a system disk. Then go to [Step 3.1](#).
 - If no, the disk is not a system disk. Then go to [Step 2.4](#).
4. Run the **df -h** command to check the system disk partition usage. Determine the role of the disk based on the disk partition name obtained in [Step 2.1](#).
5. Check whether the disk is used by HDFS or Yarn.
 - If yes, expand the disk capacity for the Core node. Then go to [Step 2.6](#).
 - If no, go to [Step 4](#).
6. Wait 2 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 3](#).

Step 3 Check whether large files are written to the disk.

1. Run the **find / -xdev -size +500M -exec ls -l {} \;** command to view files larger than 500 MB on the node. Check whether such files are written to the disk.
 - If yes, go to [Step 3.2](#).
 - If no, go to [Step 4](#).
2. Handle the large files and check whether the alarm is cleared 2 minutes later.
 - If yes, no further action is required.
 - If no, go to [Step 4](#).
3. Expand the disk capacity.
4. Wait 2 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 4](#).

Step 4 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.423 ALM-12018 Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the memory usage every 30 seconds and compares the actual memory usage with the threshold. The memory usage has a default threshold. This alarm is generated when the detected memory usage exceeds the threshold.

This alarm is cleared when the host memory usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12018	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

Memory configuration cannot meet service requirements. The memory usage reaches the upper threshold.

Procedure

Step 1 Expand the system.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host address of the alarm.

2. Log in to the node for which the alarm is generated.
3. Run `free -m | grep Mem\| | awk '{printf("%s,", ($3-$6-$7) * 100 / $2)}'` to check the system memory usage.
4. If the memory usage exceeds the threshold, expand the memory capacity.
5. Wait 5 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.424 ALM-12027 Host PID Usage Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the PID usage every 30 seconds and compares the actual PID usage with the default threshold. This alarm is generated when the PID usage exceeds the threshold.

This alarm is cleared when the host PID usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12027	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Parameter	Description
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

No PID is available for new processes and service processes are unavailable.

Possible Causes

Too many processes are running on the node. You need to increase the value of **pid_max**. The system is abnormal.

Procedure

Step 1 Increase the value of **pid_max**.

1. On the MRS cluster details page, click the alarm from the real-time alarm list. In the **Alarm Details** area, obtain the IP address of the host for which the alarm is generated.
2. Log in to the node for which the alarm is generated.
3. Run the **cat /proc/sys/kernel/pid_max** command to check the value of **pid_max**.
4. If the PID usage exceeds the threshold, open the **/etc/sysctl.conf** file and change the value of **kernel.pid_max** to twice the value of **pid_max** queried in [Step 1.3](#). If **kernel.pid_max** does not exist, add it to the end of the file.

For example, change the parameter value to **kernel.pid_max=65536** and run the following command to make the parameter take effect immediately:

```
sysctl -p
```

NOTE

The maximum value of **kernel.pid_max** is as follows:

- 32-bit OS: **32768**
 - 64-bit OS: **4194304** (22nd power of 2)
5. Wait 5 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Check whether the system environment is abnormal.

1. Contact the O&M personnel to check whether the operating system is abnormal.
 - If yes, rectify the operating system fault and go to [Step 2.2](#).
 - If no, go to [Step 3](#).
2. Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

Step 3 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.425 ALM-12028 Number of Processes in the D State on the Host Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system periodically checks the number of D state processes of user **omm** on the host every 30 seconds and compares the number with the threshold. The number of processes in the D state on the host has a default threshold. This alarm is generated when the number of processes in the D state exceeds the threshold.

This alarm is cleared when the number is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12028	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Excessive system resources are used and the service process responds slowly.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and a process is in the D state.

Procedure

Step 1 Check the process that is in the D state.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the address of the host.
2. Log in to the node for which the alarm is generated.
3. Run the following commands to switch the user:
sudo su - root
su - omm
4. Run the following command as user **omm** to view the PID of the process that is in the D state:
ps -elf | grep -v "\[thread_checkio\]" | awk 'NR!=1 {print \$2, \$3, \$4}' | grep omm | awk -F ' ' '{print \$1, \$3}' | grep D | awk '{print \$2}'
5. Check whether the command output is empty.
 - If yes, the service process is running properly. Then go to [Step 1.7](#).
 - If no, go to [Step 1.6](#).
6. Switch to user **root** and run the **reboot** command to restart the alarm host. Restarting the host brings certain risks. Ensure that the service process runs properly after the restart.
7. Wait 5 minutes and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.426 ALM-12031 User omm or Password Is About to Expire (For MRS 2.x or Earlier)

Description

The system starts at 00:00 every day to check whether user **omm** and the password are about to expire every eight hours. This alarm is generated if the user or password is about to expire in 15 days.

The alarm is cleared when the validity period of user **omm** is changed or the password is reset and the alarm handling is complete.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12031	Minor	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The node trust relationship is unavailable and Manager cannot manage the services.

Possible Causes

User **omm** or the password is about to expire.

Procedure

Step 1 Check whether user **omm** and the password in the system are valid.

1. Log in to the faulty node.
2. Run the following command to view the information about user **omm** and the password:

```
chage -l omm
```

3. Check whether the user has expired based on the system message.
 - a. View the value of **Password expires** to check whether the password is about to expire.
 - b. View the value of **Account expires** to check whether the user is about to expire.

 **NOTE**

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If yes, go to [Step 1.4](#).
 - If no, go to [Step 2](#).
4. Run the following command to modify the validity period configuration:
 - Run the following command to set a validity period for user **omm**:
chage -E 'specified date' omm
 - Run the following command to set the number of validity days for user **omm**:
chage -M 'number of days' omm
 5. Check whether the alarm is cleared automatically in the next periodic check.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.427 ALM-12032 User ommdba or Password Is About to Expire (For MRS 2.x or Earlier)

Description

The system starts at 00:00 every day to check whether user **ommdba** and the password are about to expire every eight hours. This alarm is generated if the user or password is about to expire in 15 days.

The alarm is cleared when the validity period of user **ommdba** is changed or the password is reset and the alarm handling is complete.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12032	Minor	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The OMS database cannot be managed and data cannot be accessed.

Possible Causes

User **ommdba** or the password is about to expire.

Procedure

Step 1 Check whether user **ommdba** and the password in the system are valid.

1. Log in to the faulty node.
2. Run the following command to view the information about user **ommdba** and the password:
chage -l ommdba
3. Check whether the user has expired based on the system message.
 - a. View the value of **Password expires** to check whether the password is about to expire.
 - b. View the value of **Account expires** to check whether the user is about to expire.

NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If yes, go to [Step 1.4](#).
- If no, go to [Step 2](#).

4. Run the following command to modify the validity period configuration:
 - Run the following command to set a validity period for user **ommdba**:
chage -E 'specified date' ommdba
 - Run the following command to set the number of validity days for user **ommdba**:
chage -M 'number of days' ommdba
5. Check whether the alarm is cleared automatically in the next periodic check.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.428 ALM-12033 Slow Disk Fault (For MRS 2.x or Earlier)

Description

For MRS 2.x or earlier:

- For HDDs, the alarm is triggered when any of the following conditions is met:
 - The system runs the **iostat** command every 3 seconds, and detects that the **svctm** value exceeds 1000 ms for 10 consecutive periods within 30 seconds.
 - The system runs the **iostat** command every 3 seconds, and detects that more than 60% of I/O exceeds 150 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
 - The system runs the **iostat** command every 3 seconds, and detects that the **svctm** value exceeds 1000 ms for 10 consecutive periods within 30 seconds.
 - The system runs the **iostat** command every 3 seconds, and detects that more than 60% of I/O exceeds 20 ms within 300 seconds.

This alarm is automatically cleared when the preceding conditions have not been met for 15 minutes.

For MRS 1.9.3.10 or later:

- For HDDs, the alarm is triggered when any of the following conditions is met:
 - By default, the system collects data every 3 seconds. The svctm latency reaches 1000 ms within 30 seconds in at least seven collection periods.
 - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 150 ms within 300 seconds.

- For SSDs, the alarm is triggered when any of the following conditions is met:
 - By default, the system collects data every 3 seconds. The svctm latency reaches 1000 ms within 30 seconds in at least seven collection periods.
 - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 20 ms within 300 seconds.

The collection period is 3 seconds, and the detection period is 30 or 300 seconds. This alarm is automatically cleared when none of the preceding conditions are met for three consecutive detection periods (30 or 300 seconds).

 **NOTE**

For details about how to obtain the related parameters, see [Related Information](#).

Attribute

Alarm ID	Alarm Severity	Auto Clear
12033	<ul style="list-style-type: none">• Minor: MRS 1.9.3.10 and later patch versions• Major: MRS 2.x and earlier versions	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Host Name	Specifies the host for which the alarm is generated.
DiskName	Specifies the disk for which the alarm is generated.

Impact on the System

Service performance deteriorates, service processing capabilities become poor, and services may be unavailable.

Possible Causes

The disk is aged or has bad sectors.

Procedure

Check the disk status.

- Step 1** On the MRS cluster details page, click the alarm from the real-time alarm list. In the **Alarm Details** area, obtain information about the host for which the alarm is generated and information about the faulty disk.
- Step 2** Check whether the node for which the alarm is generated is in a virtualization environment.
- If yes, go to [Step 3](#).
 - If no, go to [Step 6](#).
- Step 3** Check whether the storage performance provided by the virtualization environment meets the hardware requirements. Then, go to [Step 4](#).
- Step 4** Log in to the alarm node as user **root**, run the **df -h** command, and check whether the command output contains the value of the **DiskName** field.
- If yes, go to [Step 6](#).
 - If no, go to [Step 5](#).
- Step 5** Run the **lsblk** command to check whether the mapping between the value of **DiskName** and the disk has been created.

```
sda                8:0    0 27810G 0
├─sda1             8:1    0   509M 0 /boot
└─sda2             8:2    0 278.4G 0
   ├─system-opt (dm-0) 253:0  0   50G 0 /opt
   ├─system-root (dm-1) 253:1  0   50G 0 /
   ├─system-swap (dm-2) 253:2  0   50G 0
   └─system-var (dm-3) 253:3  0   50G 0 /var
```

- If yes, go to [Step 6](#).
 - If no, go to [Step 21](#).
- Step 6** Log in to the alarm node as user **root**, run the **lsscsi | grep "/dev/sd[x]"** command to view the disk information, and check whether RAID has been set up.

NOTE

In the command, **/dev/sd[x]** indicates the disk name obtained in [Step 1](#).

Example:

```
lsscsi | grep "/dev/sda"
```

In the command output, if **ATA**, **SATA**, or **SAS** is displayed in the third line, the disk has not been organized into a RAID group. If other information is displayed, RAID has been set up.

- If yes, go to [Step 11](#).
 - If no, go to [Step 7](#).
- Step 7** Run the **smartctl -i /dev/sd[x]** command to check whether the hardware supports the SMART tool.

Example:

```
smartctl -i /dev/sda
```

In the command output, if "SMART support is: Enabled" is displayed, the hardware supports SMART. If "Device does not support SMART" or other information is displayed, the hardware does not support SMART.

- If yes, go to [Step 8](#).
- If no, go to [Step 16](#).

Step 8 Run the `smartctl -H --all /dev/sd[x]` command to check basic SMART information and determine whether the disk is working properly.

Example:

```
smartctl -H --all /dev/sda
```

Check the value of **SMART overall-health self-assessment test result** in the command output. If the value is **FAILED**, the disk is faulty and needs to be replaced. If the value is **PASSED**, check the value of **Reallocated_Sector_Ct** or **Elements in grown defect list**. If the value is greater than 100, the disk is faulty and needs to be replaced.

- If yes, go to [Step 9](#).
- If no, go to [Step 17](#).

Step 9 Run the `smartctl -l error -H /dev/sd[x]` command to check the Glist of the disk and determine whether the disk is normal.

Example:

```
smartctl -l error -H /dev/sda
```

Check the **Command/Feattrue_name** column in the command output. If **READ SECTOR(S)** or **WRITE SECTOR(S)** is displayed, the disk has bad sectors. If other errors occur, the disk circuit board is faulty. Both errors indicate that the disk is abnormal and needs to be replaced.

If "No Errors Logged" is displayed, no error log exists. You can perform step 9 to trigger the disk SMART self-check.

- If yes, go to [Step 10](#).
- If no, go to [Step 17](#).

Step 10 Run the `smartctl -t long /dev/sd[x]` command to trigger the disk SMART self-check. After the command is executed, the time when the self-check is to be completed is displayed. After the self-check is completed, repeat [Step 8](#) and [Step 9](#) to check whether the disk is working properly.

Example:

```
smartctl -t long /dev/sda
```

- If yes, go to [Step 16](#).
- If no, go to [Step 17](#).

Step 11 Run the `smartctl -d [sat|scsi]+megaraid,[DID] -H --all /dev/sd[x]` command to check whether the hardware supports SMART.

 NOTE

- In the command, **[sat|scsi]** indicates the disk type. Both types need to be used.
- **[DID]** indicates the slot information. Slots 0 to 15 need to be used.

For example, run the following commands in sequence:

```
smartctl -d sat+megaraid,0 -H --all /dev/sda
```

```
smartctl -d sat+megaraid,1 -H --all /dev/sda
```

```
smartctl -d sat+megaraid,2 -H --all /dev/sda
```

...

Try the command combinations of different disk types and slot information. If "SMART support is: Enabled" is displayed in the command output, the disk supports SMART. Record the parameters of the disk type and slot information when a command is successfully executed. If "SMART support is: Enabled" is not displayed in the command output, the disk does not support SMART.

- If yes, go to [Step 12](#).
- If no, go to [Step 15](#).

Step 12 Run the `smartctl -d [sat|scsi]+megaraid,[DID] -H --all /dev/sd[x]` command recorded in [Step 11](#) to check basic SMART information and determine whether the disk is normal.

Example:

```
smartctl -d sat+megaraid,2 -H --all /dev/sda
```

Check the value of **SMART overall-health self-assessment test result** in the command output. If the value is **FAILED**, the disk is faulty and needs to be replaced. If the value is **PASSED**, check the value of **Reallocated_Sector_Ct** or **Elements in grown defect list**. If the value is greater than 100, the disk is faulty and needs to be replaced.

- If yes, go to [Step 13](#).
- If no, go to [Step 17](#).

Step 13 Run the `smartctl -d [sat|scsi]+megaraid,[DID] -l error -H /dev/sd[x]` command to check the Glist of the disk and determine whether the hard disk is working properly.

Example:

```
smartctl -d sat+megaraid,2 -l error -H /dev/sda
```

Check the **Command/Feattrue_name** column in the command output. If **READ SECTOR(S)** or **WRITE SECTOR(S)** is displayed, the disk has bad sectors. If other errors occur, the disk circuit board is faulty. Both errors indicate that the disk is abnormal and needs to be replaced.

If "No Errors Logged" is displayed, no error log exists. You can perform step 9 to trigger the disk SMART self-check.

- If yes, go to [Step 14](#).

- If no, go to [Step 17](#).

Step 14 Run the `smartctl -d [sat|scsi]+megaraid,[DID] -t long /dev/sd[x]` command to trigger the disk SMART self-check. After the command is executed, the time when the self-check is to be completed is displayed. After the self-check is completed, repeat [Step 12](#) and [Step 13](#) to check whether the disk is working properly.

Example:

```
smartctl -d sat+megaraid,2 -t long /dev/sda
```

- If yes, go to [Step 16](#).
- If no, go to [Step 17](#).

Step 15 If the configured RAID controller card does not support SMART, the disk does not support SMART. In this case, use the check tool provided by the corresponding RAID controller card vendor to rectify the fault. Then go to [Step 16](#).

For example, LSI is a MegaCLI tool.

Step 16 On the alarm details page, click **Clear Alarm**. Check whether the alarm is reported on the same disk again.

If the alarm is reported for more than three times, replace the disk.

- If yes, go to [Step 17](#).
- If no, no further action is required.

Replace the disk.

Step 17 On MRS Manager, choose **Alarms**.

Step 18 View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is reported.

Step 19 Replace a disk.

Step 20 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 21](#).

Collect the fault information.

Step 21 On MRS Manager, choose **System > Export Log**.

Step 22 Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Related Information

To obtain the related parameters, perform the following steps:

- For MRS 2.x or earlier versions:

Perform the following operations to detect slow disk faults:

On the Linux platform, run the **iostat -x -t 1** command to check whether the I/O is faulty. Specifically, check the **svctm** value in the red box in the figure below.

svctm indicates the I/O service time of the disk.

```
[root@opt]# iostat -x -t 1 1
Linux 4.18.0-147.5.1.6.el8.x86_64 (node-master1N3sn) 09/15/2022 _x86_64_ (4 CPU)

09/15/2022 10:57:11 AM
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           29.86    0.00   19.52    0.26    0.00   50.36

Device:            rrqm/s   wrqm/s     r/s     w/s    kB/s    kB/s avgrq-sz  avgqu-sz   await  r_await  w_await  svctm  %util
vda                 0.02    39.55    0.84   23.27   31.91   447.05    39.75    0.03    1.95    2.64    1.92    0.67   1.61
vdb                 0.01    23.61    0.21   30.88    4.08   320.62    20.88    0.01    0.86    2.08    0.85    0.71   2.21
loop0               0.00     0.00    0.00    0.00    0.01    0.00    49.94    0.00    0.31    0.31    0.00    0.29   0.00
```

- For MRS 1.9.3.10 or later patch versions:

The **svctm** value can be obtained through the following expression:

$$svctm = (tot_ticks_new - tot_ticks_old) / (rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old)$$

When the detection period is 30 seconds, if **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, then **svctm = 0**.

When the detection period is 300 seconds and **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, if **tot_ticks_new - tot_ticks_old = 0**, then **svctm = 0**; otherwise, the value of **svctm** is infinite.

The parameters in the preceding expression can be obtained as follows:

Obtain the parameter values from the data collected via the **cat /proc/diskstats** command run by the system every 3 seconds. The following shows an example.

```
omm@ ~]$ cat /proc/diskstats
253 0 vda 1101553 35446 83439787 3338546 28744856 48314024 1054257652 52667332 0 19569526 10342913 0 0 0 0
253 1 vda1 596970 25494 54533791 2565698 5446004 8749340 215777628 12114542 0 6473805 11339691 0 0 0 0
253 2 vda2 15 0 108 136 0 0 0 0 150 129 0 0 0 0
253 5 vda5 22373 1364 2525122 79502 4212374 4104759 161597984 8145606 0 3598808 6239095 0 0 0 0
253 6 vda6 11145 314 529002 85050 259201 70368 4412408 321454 0 189336 259725 0 0 0 0
253 7 vda7 157987 105 3477434 149542 6507077 1028968 140666992 14349866 0 1679035 11116587 0 0 0 0
253 8 vda8 312935 8169 22369722 458354 12179958 34360599 531802640 17724858 0 9060731 11385470 0 0 0 0
253 16 vdb 275920 21939 15977738 2171665 39472291 28236575 2653825040 482230505 0 30580346 465962048 0 0 0 0
253 17 vdb1 275439 21939 15948866 2171472 31290400 28236555 2653824832 481837775 0 30036724 465855080 0 0 0 0
253 7 loop0 356 0 17442 150 0 0 0 0 149 105 0 0 0 0
omm@ ~]$ cat /proc/diskstats
253 0 vda 1101553 35446 83439787 3338546 28747977 48319338 1054352084 52672715 0 19571460 40346640 0 0 0 0
253 1 vda1 596970 25494 54533791 2565698 5446015 8750402 215791076 12115169 0 6474429 11339985 0 0 0 0
253 2 vda2 15 0 108 136 0 0 0 0 150 129 0 0 0 0
253 5 vda5 22373 1364 2525122 79502 4212822 4105244 161614088 8146153 0 3599216 6239432 0 0 0 0
253 6 vda6 11145 314 529002 85050 259245 70433 4413368 321489 0 189389 259730 0 0 0 0
253 7 vda7 157987 105 3477434 149542 6507759 1029060 140677872 14351373 0 1679157 11117724 0 0 0 0
253 8 vda8 312935 8169 22369722 458354 12181277 34364199 531855680 17727525 0 9061647 11387424 0 0 0 0
253 16 vdb 275920 21939 15977738 2171665 39477604 28238831 2653881640 482234435 0 30581946 465964144 0 0 0 0
253 17 vdb1 275439 21939 15948866 2171472 31293358 28238811 2653881432 481841639 0 30038274 465857164 0 0 0 0
253 7 loop0 356 0 17442 150 0 0 0 0 149 105 0 0 0 0
```

In the data collected for the first time, the number in the fourth column is the value of **rd_ios_old**, the number in the eighth column is the value of **wr_ios_old**, and the number in the thirteenth column is the value of **tot_ticks_old**.

In the data collected for the second time, the number in the fourth column is the value of **rd_ios_new**, the number in the eighth column is the value of **wr_ios_new**, and the number in the thirteenth column is the value of **tot_ticks_new**.

In this case, the value of **svctm** is as follows:

$$(19571460 - 19569526) / (1101553 + 28747977 - 1101553 - 28744856) = 0.6197$$

7.12.429 ALM-12034 Periodic Backup Failure (For MRS 2.x or Earlier)

Description

This alarm is generated when a periodic backup task fails to be executed. This alarm is cleared when the next backup task is executed successfully.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12034	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TaskName	Specifies the task name.

Impact on the System


No backup package is available for a long time, so the system cannot be restored in case of exceptions.

Possible Causes

The alarm cause depends on the task details. Handle the alarm according to the logs and alarm details.

Procedure

Checking whether the disk space is insufficient

- Step 1** On MRS Manager, choose **Alarms**.
- Step 2** In the alarm list, click  of the alarm and obtain the task name from the **Location** area.
- Step 3** Choose **System > Back Up Data**.

- Step 4** Search for the backup task based on the task name and choose **More > View History** in the **Operation** column to view detailed information about the backup task.
- Step 5** Choose **Details > View** and check whether message "Failed to backup xx due to insufficient disk space, move the data in the /srv/BigData/LocalBackup directory to other directories." exists.
- If yes, go to [Step 6](#).
 - If no, go to [Step 13](#).
- Step 6** Choose **Backup Path > View** to obtain the backup path.
- Step 7** Log in to the node as user **root** and view the mounting details of the node.
- df -h**
- Step 8** Check whether the available space of the node to which the backup path is mounted is less than 20 GB.
- If yes, go to [Step 9](#).
 - If no, go to [Step 13](#).
- Step 9** Check whether the backup package exists in the backup directory and whether the available space of the node to which the backup directory is mounted is less than 20 GB.
- If yes, go to [Step 10](#).
 - If no, go to [Step 13](#).
- Step 10** Ensure that the available space of the node to which the backup directory is mounted to be greater than 20 GB by moving backup packages out of the backup directory or deleting the backup packages.
- Step 11** Start the backup task again and check whether the backup task is executed.
- If yes, go to [Step 12](#).
 - If no, go to [Step 13](#).
- Step 12** After 2 minutes, check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 13](#).
- Collecting fault information**
- Step 13** On MRS Manager, choose **System > Export Log**.
- Step 14** Contact the O&M engineers and send the collected logs.
- End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Reference

None

7.12.430 ALM-12035 Unknown Data Status After Recovery Task Failure (For MRS 2.x or Earlier)

Description

If a recovery task fails, the system attempts to automatically roll back. If the rollback fails, data may be lost. If this occurs, an alarm is reported. This alarm is cleared when the recovery task is successfully executed later.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12035	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TaskName	Specifies the task name.

Impact on the System

The data may be lost or the data status may be unknown, which may affect services.

Possible Causes

The possible cause of this alarm is that the component status does not meet the requirements before the restoration task is executed or an error occurs in a step during the restoration task. The error depends on the task details. You can obtain logs and task details to handle the alarm.

Procedure

Checking the component status

- Step 1** Log in to MRS Manager and choose **Services**. On the page that is displayed, check whether the running status of the components meets the requirements. (OMS and DBService must be in the normal status, and other components must be stopped.)

- If yes, go to [Step 7](#).
- If no, go to [Step 2](#).

Step 2 Restore the component status as required and start the recovery task again.

Step 3 Log in to MRS Manager and choose **Alarms**. In the alarm list, click the row containing the alarm and obtain the task name from the **Location** area.

Step 4 Choose **System > Recovery Management**. Search for the restoration task based on the task name and view the task details.

Step 5 Start the restoration task and check whether the task is executed.

- If yes, go to [Step 6](#).
- If no, go to [Step 7](#).

Step 6 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 7](#).

Collecting fault information

Step 7 On MRS Manager, choose **System > Export Log**.

Step 8 Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Reference

None

7.12.431 ALM-12037 NTP Server Abnormal (For MRS 2.x or Earlier)

Description

This alarm is generated when the NTP server is abnormal.

This alarm is cleared when the NTP server recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12037	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the IP address of the NTP server for which the alarm is generated.

Impact on the System

The NTP server configured on the active OMS node is abnormal. In this case, the active OMS node cannot synchronize time with the NTP server and a time offset may be generated in the cluster.

Possible Causes

- The NTP server network is faulty.
- The NTP server authentication fails.
- The time cannot be obtained from the NTP server.
- The time obtained from the NTP server is not continuously updated.

Procedure

Step 1 Check the NTP server network.

1. On the MRS cluster details page, click the alarm from the real-time alarm list.
2. In the **Alarm Details** area, view the additional information to check whether the NTP server fails to be pinged.
 - If yes, go to [Step 1.3](#).
 - If no, go to [Step 2](#).
3. Contact the O&M personnel to check the network configuration and ensure that the network between the NTP server and the active OMS node is in normal state. Then, check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Check whether the NTP server authentication fails.

1. Log in to the active management node.
2. Run **ntpq -np** to check whether the NTP server authentication fails. If **refid** of the NTP server is **.AUTH.**, the authentication fails.
 - If yes, go to [Step 5](#).
 - If no, go to [Step 3](#).

Step 3 Check whether the time can be obtained from the NTP server.

1. View the alarm additional information to check whether the time cannot be obtained from the NTP server.
 - If yes, go to [Step 3.2](#).
 - If no, go to [Step 4](#).
2. Contact the O&M personnel to rectify the NTP server fault. After the NTP server is in normal state, check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 4](#).

Step 4 Check whether the time obtained from the NTP server fails to be updated.

1. View the alarm additional information to check whether the time obtained from the NTP server fails to be updated.
 - If yes, go to [Step 4.2](#).
 - If no, go to [Step 5](#).
2. Contact the provider of the NTP server to rectify the NTP server fault. After the NTP server is in normal state, check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 5](#).

Step 5 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.432 ALM-12038 Monitoring Indicator Dump Failure (For MRS 2.x or Earlier)

Description

This alarm is generated when dumping fails after monitoring indicator dumping is configured on MRS Manager.

This alarm is cleared when dumping is successful.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12038	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The upper-layer management system fails to obtain monitoring indicators from the MRS Manager system.

Possible Causes

- The server cannot be connected.
- The save path on the server cannot be accessed.
- The monitoring indicator file fails to be uploaded.

Procedure

- Step 1** Contact the O&M personnel to check whether the network connection between the MRS Manager system and the server is normal.
 - If yes, go to [Step 3](#).
 - If no, go to [Step 2](#).
- Step 2** Contact the O&M personnel to restore the network and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 3](#).
- Step 3** Choose **System > Monitor Dumping Configuration** and check whether the FTP username, password, port, dump mode, and public key configured on the monitoring indicator dumping configuration page are consistent with those on the server.
 - If yes, go to [Step 5](#).
 - If no, go to [Step 4](#).
- Step 4** Enter the correct configuration, click **OK**, and check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 5](#).
- Step 5** Choose **System > Monitor Dumping Configuration** and check the configuration items, including the FTP username, save path, and dumping mode.
 - If the FTP mode is used, go to [Step 6](#).

- If the SFTP mode is used, go to [Step 7](#).
- Step 6** Log in to the server. In the default path, check whether the save path (relative path) has the read and write permission on the FTP username.
- If yes, go to [Step 9](#).
 - If no, go to [Step 8](#).
- Step 7** Log in to the server. In the default path, check whether the save path (absolute path) has the read and write permission on the FTP username.
- If yes, go to [Step 9](#).
 - If no, go to [Step 8](#).
- Step 8** Add the read and write permission and check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 9](#).
- Step 9** Log in to the server and check whether the save path has sufficient disk space.
- If yes, go to [Step 11](#).
 - If no, go to [Step 10](#).
- Step 10** Delete unnecessary files or go to the monitoring indicator dumping configuration page to change the save path. Check whether the alarm is cleared.
- If yes, no further action is required.
 - If no, go to [Step 11](#).
- Step 11** Collect fault information.
1. On MRS Manager, choose **System** > **Export Log**.
 2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.433 ALM-12039 GaussDB Data Is Not Synchronized (For MRS 2.x or Earlier)

Description

The system checks the data synchronization status between the active and standby GaussDB nodes every 10 seconds. This alarm is generated when the synchronization status cannot be queried for six consecutive times or when the synchronization status is abnormal.

This alarm is cleared when the data synchronization status is normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12039	Critical	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.
Local GaussDB HA IP	HA IP address of the local GaussDB.
Peer GaussDB HA IP	HA IP address of the peer GaussDB.
SYNC_PERCENT	Synchronization percentage.

Impact on the System

When data is not synchronized between the active and standby GaussDBs, the data may be lost or abnormal if the active instance becomes abnormal.

Possible Causes

- The network between the active and standby nodes is unstable.
- The standby GaussDB is abnormal.
- The disk space of the standby node is full.

Procedure

Step 1 Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the IP address of the standby GaussDB node.

Step 2 Log in to the active management node.

Step 3 Run the following command to check whether the standby GaussDB is reachable:

ping *heartbeat IP address of the standby GaussDB*

If yes, go to [Step 6](#).

If no, go to [Step 4](#).

Step 4 Contact the O&M personnel to check whether the network is faulty.

- If yes, go to [Step 5](#).

- If no, go to [Step 6](#).

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to [Step 6](#).

Step 6 Log in to the standby GaussDB node.

Step 7 Run the following commands to switch the user:

```
sudo su - root
```

```
su - omm
```

Step 8 Go to the `${BIGDATA_HOME}/om-0.0.1/sbin/` directory.

Run the following command to check whether the resource status of the standby GaussDB is normal:

```
sh status-oms.sh
```

In the command output, check whether the following information is displayed in the row where **ResName** is **gaussDB**:

```
10_10_10_231 gaussDB Standby_normal Normal Active_standby
```

- If yes, go to [Step 9](#).
- If no, go to [Step 15](#).

Step 9 Log in to the standby GaussDB node.

Step 10 Run the following commands to switch the user:

```
sudo su - root
```

```
su - omm
```

Step 11 Run the `echo ${BIGDATA_DATA_HOME}/dbdata_om` command to obtain the GaussDB data directory.

Step 12 Run the `df -h` command to check the system disk partition usage.

Step 13 Check whether the disk where the GaussDB data directory is mounted is full.

- If yes, go to [Step 14](#).
- If no, go to [Step 15](#).

Step 14 Contact the O&M personnel to expand the disk capacity. After capacity expansion, wait 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 15](#).

Step 15 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.434 ALM-12040 Insufficient System Entropy (For MRS 2.x or Earlier)

Description

The system checks the entropy at 00:00:00 every day and performs five consecutive checks each time. First, the system checks whether the rng-tools tool is enabled and correctly configured. If not, the system checks the current entropy. This alarm is generated if the entropy is less than 500 in the five checks.

This alarm is cleared if the true random number mode is configured, random numbers are configured in pseudo-random number mode, or neither the true random number mode nor the pseudo-random number mode is configured but the entropy is greater than or equal to 500 in at least one check among the five checks.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12040	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Decryption failures occur and functions related to decryption are affected, for example, DBService installation.

Possible Causes

The rngd service is abnormal.

Procedure

- Step 1** Go to the cluster details page and choose **Alarms**.
- Step 2** View the alarm details to obtain the value of the **HostName** field in **Location**.
- Step 3** Log in to the node for which the alarm is generated and run the **sudo su - root** command to switch to user **root**.
- Step 4** Run the **/bin/rpm -qa | grep -w "rng-tools"** command. If the command is executed successfully, run the **ps -ef | grep -v "grep" | grep rngd | tr -d " " | grep "\-o/dev/random" | grep "\-r/dev/urandom"** command and view the command output.
- If the command is executed successfully, the rngd service is installed, correctly configured, and is running properly. Go to [Step 6](#).
 - If the command is not executed successfully, the rngd service is not running properly. Then go to [Step 5](#).
- Step 5** Run the following command to start the rngd service:
- ```
echo 'EXTRAOPTIONS="-r /dev/urandom -o /dev/random"' >> /etc/sysconfig/rngd

service rngd start
```
- Step 6** Run the **service rngd status** command to check whether the rngd service is in the running state.
- If yes, go to [Step 7](#).
  - If no, go to [Step 8](#).
- Step 7** Wait until 00:00:00 when the system checks the entropy again. Check whether the alarm is cleared automatically.
- If yes, no further action is required.
  - If no, go to [Step 8](#).
- Step 8** Collect fault information.
1. On MRS Manager, choose **System > Export Log**.
  2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.435 ALM-12041 Permission of Key Files Is Abnormal (For MRS 2.x or Earlier)

### Description

The system checks the permission, users, and user groups of key directories or files every hour. This alarm is generated if any of these is abnormal.



This alarm is cleared after the problem that causes abnormal permission, users, or user groups is solved.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12041    | Major          | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |
| PathName    | Specifies the file path or file name.                   |

## Impact on the System

System functions are unavailable.

## Possible Causes

The user has manually modified the file permission, user information, or user groups, or the system has experienced an unexpected power-off.

## Procedure

### Step 1 Check the file permission.

1. Go to the MRS cluster details page and choose **Alarms**.
2. In the details of the alarm, query the **HostName** (name of the alarmed host) and **PathName** (path or name of the involved file).
3. Log in to the alarmed node.
4. Run the **ll PathName** command to query the current user, permission, and user group of the file or path.
5. Go to the **#{BIGDATA\_HOME}/nodeagent/etc/agent/autocheck** directory and run the **vi keyfile** command. Search for the name of the involved file and query the correct permission of the file.
6. Compare the actual permission of the file with the permission obtained in [Step 1.5](#). If they are different, change the actual permission, user information, and user group to the correct values.

7. Wait until the next system check is complete and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.436 ALM-12042 Key File Configurations Are Abnormal (For MRS 2.x or Earlier)

### Description

The system checks key file configurations every hour. This alarm is generated if any key configuration is abnormal.

This alarm is cleared after the configuration becomes normal.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12042    | Major          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |
| PathName    | Specifies the file path or file name.                   |

### Impact on the System

Functions related to the file are abnormal.

## Possible Causes

The user has manually modified the file configurations or the system has experienced an unexpected power-off.

## Procedure

**Step 1** Check the file configurations.

1. Go to the MRS cluster details page and choose **Alarms**.
2. In the details of the alarm, query the **HostName** (name of the alarmed host) and **PathName** (path or name of the involved file).
3. Log in to the alarmed node.
4. Manually check and modify the file configurations according to the criteria in [Related Information](#).
5. Wait until the next system check is complete and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

- **Checking /etc/fstab**

Check whether partitions configured in **/etc/fstab** exist in **/proc/mounts** and whether swap partitions configured in **/etc/fstab** match those in **/proc/swaps**.

- **Checking /etc/hosts**

Run the **cat /etc/hosts** command. If any of the following situations exists, the file configurations are abnormal.

- The **/etc/hosts** file does not exist.
- The host name is not configured in the file.
- The IP address of the host is duplicate.
- The IP address of the host does not exist in the **ipconfig** list.
- An IP address in the file is used by multiple hosts.

## 7.12.437 ALM-12043 DNS Parsing Duration Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the DNS parsing duration every 30 seconds. This alarm is generated when the DNS parsing duration exceeds the threshold (the default threshold is 20,000 ms) for multiple times (the default value is 2).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Network Status > DNS Resolution Duration > DNS Resolution Duration**.

This alarm is cleared when **hit number** is **1** and the DNS resolution duration is less than or equal to the threshold. This alarm is cleared when **hit number** is not **1** and the DNS resolution duration is less than or equal to 90% of the threshold.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12043    | Major          | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

- Kerberos-based secondary authentication is slow.
- The ZooKeeper service is abnormal.
- The node is faulty.

## Possible Causes

- The node is configured with the DNS client.
- The node is equipped with the DNS server and the DNS server is started.

## Procedure

**Check whether the node is configured with the DNS client.**

- Step 1** Go to the MRS cluster details page and choose **Alarms**.
- Step 2** View the alarm details to obtain the value of the **HostName** field in **Location**.
- Step 3** Use PuTTY to log in to the node for which the alarm is generated as user **root**.
- Step 4** Run the **cat /etc/resolv.conf** command to check whether the DNS client is installed.

If information similar to the following is displayed, the DNS client is installed and started:

```
nameserver 10.2.3.4
nameserver 10.2.3.4
```

- If yes, go to [Step 5](#).
- If no, go to [Step 7](#).

**Step 5** Run the `vi /etc/resolv.conf` command to comment out the following content using the number signs (#) and save the file:

```
nameserver 10.2.3.4
nameserver 10.2.3.4
```

**Step 6** Check whether this alarm is cleared after 5 minutes.

- If yes, no further action is required.
- If no, go to [Step 7](#).

**Check whether the node is equipped with the DNS server and the DNS server is started.**

**Step 7** Run the `service named status` command to check whether the DNS service is installed on the node.

If information similar to the following is displayed, the DNS server is installed and started:

```
Checking for nameserver BIND
version: 9.6-ESV-R7-P4
CPUs found: 8
worker threads: 8
number of zones: 17
debug level: 0
xfers running: 0
xfers deferred: 0
soa queries in progress: 0
query logging is ON
recursive clients: 4/0/1000
tcp clients: 0/100
server is up and running
```

- If yes, go to [Step 8](#).
- If no, go to [Step 10](#).

**Step 8** Run the `service named stop` command to stop the DNS server.

**Step 9** Check whether this alarm is cleared after 5 minutes.

- If yes, no further action is required.
- If no, go to [Step 10](#).

**Step 10** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.438 ALM-12045 Read Packet Dropped Rate Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the read packet dropped rate every 30 seconds. This alarm is generated when the read packet dropped rate exceeds the threshold (the default threshold is 0.5%) for multiple times (the default value is 5).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Network Reading > Network Read Packet Rate Information > Read Packet Dropped Rate**.

This alarm is cleared when **hit number** is 1 and the read packet dropped rate is less than or equal to the threshold. This alarm is cleared when **hit number** is greater than 1 and the read packet dropped rate is less than or equal to 90% of the threshold.

The alarm detection is disabled by default. If you want to enable this function, check whether this function can be enabled based on Checking System Environments.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12045    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                  |
|-------------------|--------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.      |
| RoleName          | Specifies the role for which the alarm is generated.         |
| HostName          | Specifies the host for which the alarm is generated.         |
| NetworkCardName   | Specifies the network port for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm.            |

### Impact on the System

The service performance deteriorates or some services time out.

Risk warning: In SUSE kernel 3.0 or later or Red Hat 7.2, the system kernel modifies the mechanism for counting the number of dropped read packets. In this

case, this alarm may be generated even if the network is running properly, but services are not affected. You are advised to check the system environment first.

## Possible Causes

- An OS exception occurs.
- The NICs are bonded in active/standby mode.
- The alarm threshold is improperly configured.
- The network environment is abnormal.

## Procedure

### View the network packet dropped rate.

- Step 1** Use PuTTY to log in to any non-alarm node in the cluster as user **omm** and run the **ping IP address of the node for which the alarm is generated -c 100** command to check whether packet drop occurs on the network.

```
ping 10.10.10.12 -c 5
PING 10.10.10.12 (10.10.10.12) 56(84) bytes of data.
64 bytes from 10.10.10.11: icmp_seq=1 ttl=64 time=0.033 ms
64 bytes from 10.10.10.11: icmp_seq=2 ttl=64 time=0.034 ms
64 bytes from 10.10.10.11: icmp_seq=3 ttl=64 time=0.021 ms
64 bytes from 10.10.10.11: icmp_seq=4 ttl=64 time=0.033 ms
64 bytes from 10.10.10.11: icmp_seq=5 ttl=64 time=0.030 ms
--- 10.10.10.12 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4001ms rtt min/avg/max/mdev =
0.021/0.030/0.034/0.006 ms
```

#### NOTE

- *IP address of the node for which the alarm is generated.* Query the IP address of the node for which the alarm is generated on the node management page of the MRS cluster details page based on the value of **HostName** in the alarm location information. Check both the IP addresses of the management plane and service plane.
- **-c**: number of check times. The default value is **100**.
- If yes, go to [Step 11](#).
- If no, go to [Step 2](#).

### Check the system environment.

- Step 2** Use PuTTY to log in to the active OMS node or the node for which the alarm is generated as user **omm**.

- Step 3** Run the **cat /etc/\*-release** command to check the OS type.

- If the OS is EulerOS, go to [Step 4](#).

```
cat /etc/*-release
EulerOS release 2.0 (SP2)
```
- If the OS is SUSE, go to [Step 5](#).

```
cat /etc/*-release
SUSE Linux Enterprise Server 11 (x86_64)
VERSION = 11
PATCHLEVEL = 3
```
- Otherwise, go to [Step 11](#).

- Step 4** Run the **cat /etc/euleros-release** command to check whether the OS version is EulerOS 2.2.

```
cat/etc/euleros-release
EulerOS release 2.0 (SP2)
```

- If yes, the alarm sending function cannot be enabled. Go to [Step 6](#).
- If no, go to [Step 11](#).

**Step 5** Run the `cat /proc/version` command to check whether the SUSE kernel version is 3.0 or later.

```
cat /proc/version
Linux version 3.0.101-63-default (geeko@buildhost) (gcc version 4.3.4 [gcc-4_3-branch revision 152973]
(SUSE Linux)) #1 SMP Tue Jun 23 16:02:31 UTC 2015 (4b89d0c)
```

- If yes, the alarm sending function cannot be enabled. Go to [Step 6](#).
- If no, go to [Step 11](#).

**Step 6** Log in to MRS Manager and choose **System > Configuration > Threshold Configuration**.

**Step 7** In the navigation pane of the **Threshold Configuration** page, choose **Network Reading > Network Read Packet Rate Information > Read Packet Dropped Rate**. In the right pane, check whether **Send Alarm** is selected.

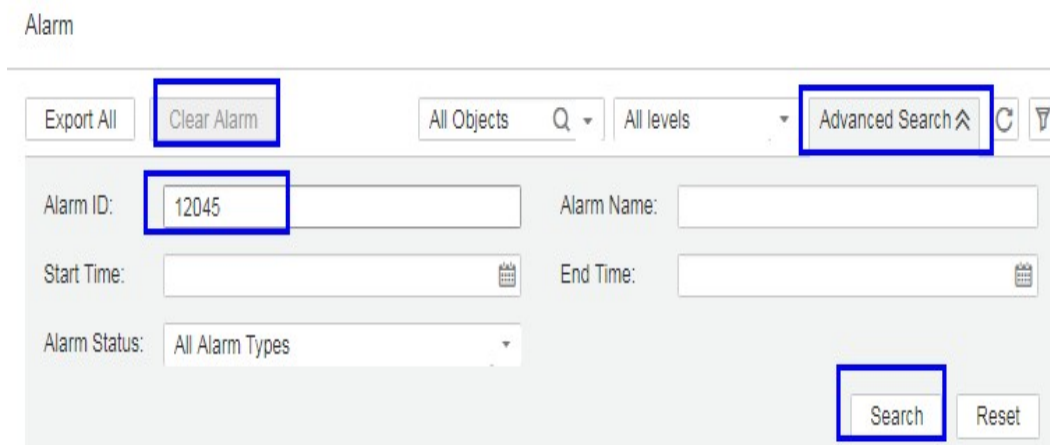
- If yes, the alarm sending function is enabled. Go to [Step 8](#).
- If no, the alarm sending function is disabled. Go to [Step 10](#).

**Step 8** In the right pane, deselect **Send Alarm** to shield alarm "Network Read Packet Dropped Rate Exceeds the Threshold."

**Step 9** Go to the MRS cluster details page and choose **Alarms**.

**Step 10** Search for alarm 12045 and manually clear the alarms that are not automatically cleared. No further action is required.

**Figure 7-250** Alarm Management



**NOTE**

The ID of alarm Network Read Packet Dropped Rate Exceeds the Threshold is 12045.

**Check whether the NICs are bonded in active/standby mode.**

**Step 11** Use PuTTY to log in to the node for which the alarm is generated as user **omm** and run the `ls -l /proc/net/bonding` command to check whether the `/proc/net/bonding` directory exists on the node.

- If yes, as shown in the following figure, the bond mode is configured for the node. Go to [Step 12](#).



```
ls -l /proc/net/bonding/
total 0
-r--r--r-- 1 root root 0 Oct 11 17:35 bond0
```

- If no, the bond mode is not configured for the node. Go to [Step 14](#).

```
ls -l /proc/net/bonding/
ls: cannot access /proc/net/bonding/: No such file or directory
```

**Step 12** Run the `cat /proc/net/bonding/bond0` command to check whether the value of **Bonding Mode** in the configuration file is **fault-tolerance**.

#### NOTE

In the preceding command, **bond0** is the name of the bond configuration file. Use the file name obtained in [Step 11](#).

```
cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
```

```
Bonding Mode: fault-tolerance (active-backup)
Primary Slave: eth1 (primary_reselect always)
Currently Active Slave: eth1
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
```

```
Slave Interface: eth0
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 1
Slave queue ID: 0
```

```
Slave Interface: eth1
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 1
Slave queue ID: 0
```

- If yes, the NICs are bonded in active/standby mode. Go to [Step 13](#).
- If no, go to [Step 14](#).

**Step 13** Check whether the NIC specified by **NetworkCardName** in the alarm details is the standby NIC.

- If yes, the alarm of the standby NIC cannot be automatically cleared. Manually clear the alarm on the alarm management page. No further action is required.
- If no, go to [Step 14](#).

#### NOTE

To determine the standby NIC, check the `/proc/net/bonding/bond0` configuration file. If the NIC name corresponding to **NetworkCardName** is **Slave Interface** but not **Currently Active Slave** (the current active NIC), the NIC is the standby one.

#### Check whether the threshold is set properly.

**Step 14** Log in to MRS Manager and check whether the threshold (configurable, 0.5% by default) is appropriate.

- If yes, go to [Step 17](#).
- If no, go to [Step 15](#).

**Step 15** Choose **System > Threshold Configuration > Device > Host > Network Reading > Network Read Packet Rate Information > Read Packet Dropped Rate** and change the alarm threshold based on the actual service usage.

**Step 16** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 17](#).

**Check whether the network is normal.**

**Step 17** Contact the system administrator to check whether the network is normal.

- If yes, rectify the network fault and go to [Step 18](#).
- If no, go to [Step 19](#).

**Step 18** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 19](#).

**Step 19** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.439 ALM-12046 Write Packet Dropped Rate Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the write packet dropped rate every 30 seconds. This alarm is generated when the write packet dropped rate exceeds the threshold (the default threshold is 0.5%) for multiple times (the default value is 5).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Network Writing > Network Write Packet Rate Information > Write Packet Dropped Rate**.

When the **hit number** is **1**, this alarm is cleared when the network write packet dropped rate is less than or equal to the threshold. When the **hit number** is greater than **1**, this alarm is cleared when the network write packet dropped rate is less than or equal to 90% of the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12046    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                  |
|-------------------|--------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.      |
| RoleName          | Specifies the role for which the alarm is generated.         |
| HostName          | Specifies the host for which the alarm is generated.         |
| NetworkCardName   | Specifies the network port for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm.            |

## Impact on the System

The service performance deteriorates or some services time out.

## Possible Causes

- The alarm threshold is improperly configured.
- The network environment is abnormal.

## Procedure

**Check whether the threshold is set properly.**

**Step 1** Log in to MRS Manager and check whether the threshold (configurable, 0.5% by default) is appropriate.

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

**Step 2** Choose **System > Threshold Configuration > Device > Host > Network Write Information > Network Write Packet Rate > Write Packet Dropped Rate** and change the alarm threshold based on the actual service usage.

**Step 3** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the network is normal.**

**Step 4** Contact the system administrator to check whether the network is normal.

- If yes, rectify the network fault and go to [Step 5](#).

- If no, go to [Step 6](#).

**Step 5** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Step 6** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.440 ALM-12047 Read Packet Error Rate Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the read packet error rate every 30 seconds. This alarm is generated when the read packet error rate exceeds the threshold (the default threshold is **0.5%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Network Reading > Network Read Packet Rate Information > Read Packet Error Rate**.

If the **hit number** is **1**, this alarm is cleared when the read packet error rate is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the read packet error rate is less than or equal to 90% of the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12047    | Major          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |

| Parameter         | Description                                                  |
|-------------------|--------------------------------------------------------------|
| HostName          | Specifies the host for which the alarm is generated.         |
| NetworkCardName   | Specifies the network port for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm.            |

## Impact on the System

The communication is intermittently interrupted, and services time out.

## Possible Causes

- The alarm threshold is improperly configured.
- The network environment is abnormal.

## Procedure

**Check whether the threshold is set properly.**

- Step 1** Log in to MRS Manager and check whether the threshold (configurable, 0.5% by default) is appropriate.
- If yes, go to [Step 4](#).
  - If no, go to [Step 2](#).

- Step 2** Choose **System > Threshold Configuration > Device > Host > Network Reading > Network Read Packet Rate Information > Read Packet Error Rate** and change the alarm threshold based on the actual service usage.

- Step 3** Wait 5 minutes and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 4](#).

**Check whether the network is normal.**

- Step 4** Contact the system administrator to check whether the network is normal.
- If yes, rectify the network fault and go to [Step 5](#).
  - If no, go to [Step 6](#).

- Step 5** Wait 5 minutes and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 6](#).

- Step 6** Collect fault information.
1. On MRS Manager, choose **System > Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.441 ALM-12048 Write Packet Error Rate Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the write packet error rate every 30 seconds. This alarm is generated when the write packet error rate exceeds the threshold (the default threshold is **0.5%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Network Writing > Network Write Packet Rate Information > Write Packet Error Rate**.

If **hit number** is **1**, this alarm is cleared when the write packet error rate is less than or equal to the threshold. If **hit number** is greater than **1**, this alarm is cleared when the write packet error rate is less than or equal to 90% of the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12048    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                  |
|-------------------|--------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.      |
| RoleName          | Specifies the role for which the alarm is generated.         |
| HostName          | Specifies the host for which the alarm is generated.         |
| NetworkCardName   | Specifies the network port for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm.            |

## Impact on the System

The communication is intermittently interrupted, and services time out.

## Possible Causes

- The alarm threshold is improperly configured.
- The network environment is abnormal.

## Procedure

**Check whether the threshold is set properly.**

**Step 1** Log in to MRS Manager and check whether the threshold (configurable, 0.5% by default) is appropriate.

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

**Step 2** Choose **System > Threshold Configuration > Device > Host > Network Writing > Network Write Packet Rate Information > Write Packet Error Rate** and change the alarm threshold based on the actual service usage.

**Step 3** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Check whether the network is normal.**

**Step 4** Contact the system administrator to check whether the network is normal.

- If yes, rectify the network fault and go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Step 6** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.442 ALM-12049 Read Throughput Rate Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the read throughput rate every 30 seconds. This alarm is generated when the read throughput rate exceeds the threshold (the default threshold is **80%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Network Reading > Network Read Throughput Rate > Read Throughput Rate**.

If the **hit number** is **1**, this alarm is cleared when the read throughput rate is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the read throughput rate is less than or equal to 90% of the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12049    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                  |
|-------------------|--------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.      |
| RoleName          | Specifies the role for which the alarm is generated.         |
| HostName          | Specifies the host for which the alarm is generated.         |
| NetworkCardName   | Specifies the network port for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm.            |

### Impact on the System

The service system runs abnormally or is unavailable.

### Possible Causes

- The alarm threshold is improperly configured.



- The network port rate does not meet service requirements.

## Procedure

### Check whether the threshold is set properly.

- Step 1** Log in to MRS Manager and check whether the threshold (configurable, 80% by default) is appropriate.
- If yes, go to [Step 2](#).
  - If no, go to [Step 4](#).
- Step 2** Choose **System > Threshold Configuration > Device > Host > Network Reading > Network Read Throughput Rate > Read Throughput Rate** to change the alarm threshold based on the actual service usage.
- Step 3** Wait 5 minutes and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 4](#).

### Check whether the network port rate meets the requirements.

- Step 4** In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address and network port name of the host for which the alarm is generated.
- Step 5** Use PuTTY to log in to the host for which the alarm is generated as user **root**.
- Step 6** Run the **ethtool *network port name*** command to check the maximum network port rate **Speed**.

#### NOTE

In a VM environment, you may fail to obtain the network port rate by running commands. You are advised to contact the system administrator to check whether the network port rate meets the requirements.

- Step 7** If the read throughput rate exceeds the threshold, contact the system administrator to increase the network port rate.
- Step 8** Check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 9](#).
- Step 9** Collect fault information.
1. On MRS Manager, choose **System > Export Log**.
  2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.443 ALM-12050 Write Throughput Rate Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the write throughput rate every 30 seconds. This alarm is generated when the write throughput rate exceeds the threshold (the default threshold is **80%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Network Writing > Network Write Throughput Rate > Write Throughput Rate**.

If the **hit number** is **1**, this alarm is cleared when the write throughput rate is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the write throughput rate is less than or equal to 90% of the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12050    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                  |
|-------------------|--------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.      |
| RoleName          | Specifies the role for which the alarm is generated.         |
| HostName          | Specifies the host for which the alarm is generated.         |
| NetworkCardName   | Specifies the network port for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm.            |

### Impact on the System

The service system runs abnormally or is unavailable.

### Possible Causes

- The alarm threshold is improperly configured.

- The network port rate does not meet service requirements.

## Procedure

### Check whether the threshold is set properly.

**Step 1** Log in to MRS Manager and check whether the threshold (configurable, 80% by default) is appropriate.

- If yes, go to [Step 4](#).
- If no, go to [Step 2](#).

**Step 2** Choose **System > Threshold Configuration > Device > Host > Network Writing > Network Write Throughput Rate > Write Throughput Rate** to change the alarm threshold based on the actual service usage.

**Step 3** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

### Check whether the network port rate meets the requirements.

**Step 4** In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address and network port of the host for which the alarm is generated.

**Step 5** Use PuTTY to log in to the host for which the alarm is generated as user **root**.

**Step 6** Run the **ethtool network port name** command to check the maximum network port rate **Speed**.

#### NOTE

In a VM environment, you may fail to obtain the network port rate by running commands. You are advised to contact the system administrator to check whether the network port rate meets the requirements.

**Step 7** If the write throughput rate exceeds the threshold, contact the system administrator to increase the network port rate.

**Step 8** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Step 9** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.444 ALM-12051 Disk Inode Usage Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the disk inode usage every 30 seconds. This alarm is generated when the disk inode usage exceeds the threshold (the default threshold is 80%) for multiple times (the default value is 5).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Disk > Disk Inode Usage > Disk Inode Usage**.

If the **hit number** is **1**, this alarm is cleared when the disk inode usage is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the disk inode usage is less than or equal to 90% of the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12051    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                    |
|-------------------|----------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.        |
| RoleName          | Specifies the role for which the alarm is generated.           |
| HostName          | Specifies the host for which the alarm is generated.           |
| PartitionName     | Specifies the disk partition for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm.              |

### Impact on the System

Data cannot be written to the file system.

### Possible Causes

- There are too many small files on the disk.
- The system is abnormal.

## Procedure

### There are too many small files on the disk.

- Step 1** Go to the MRS cluster details page and choose **Alarms**.
- Step 2** In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address and disk partitions of the host for which the alarm is generated.
- Step 3** Use PuTTY to log in to the host for which the alarm is generated as user **root**.
- Step 4** Run the **df -i *partition name*** command to check the current inode usage of the disk.
- Step 5** If the inode usage exceeds the threshold, manually check whether the small files in the partition can be deleted.
- If yes, delete the files and go to **Step 6**.
  - If no, adjust the capacity. Then go to **Step 7**.
- Step 6** Wait 5 minutes and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 7**.
- Check whether the system environment is normal.**
- Step 7** Contact the operating system maintenance personnel to check whether the system environment is abnormal.
- If yes, rectify the operating system fault and go to **Step 8**.
  - If no, go to **Step 9**.
- Step 8** Wait 5 minutes and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to **Step 9**.
- Step 9** Collect fault information.
1. On MRS Manager, choose **System > Export Log**.
  2. Contact the O&M engineers and send the collected logs.
- End

## Reference

None

## 7.12.445 ALM-12052 Usage of Temporary TCP Ports Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the usage of temporary TCP ports every 30 seconds. This alarm is generated when the usage of temporary TCP ports exceeds the threshold (the default threshold is **80%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System > Threshold Configuration > Host > Network Status > TCP Ephemeral Port Usage > TCP Ephemeral Port Usage**.

If the **hit number** is **1**, this alarm is cleared when the usage of temporary TCP ports is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the usage of temporary TCP ports is less than or equal to 90% of the threshold.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12052    | Major          | Yes        |

## Parameters

| Parameter         | Description                                             |
|-------------------|---------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated. |
| RoleName          | Specifies the role for which the alarm is generated.    |
| HostName          | Specifies the host for which the alarm is generated.    |
| Trigger Condition | Specifies the threshold for triggering the alarm.       |

## Impact on the System

Services on the host fail to establish connections with the external and services are interrupted.

## Possible Causes

- The temporary ports do not meet service requirements.
- The system is abnormal.

## Procedure

### Expand the range of temporary ports.

- Step 1** Go to the MRS cluster details page and choose **Alarms**.
- Step 2** In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address of the host for which the alarm is generated.
- Step 3** Use PuTTY to log in to the host for which the alarm is generated as user **omm**.

**Step 4** Run the `cat /proc/sys/net/ipv4/ip_local_port_range |cut -f 1` command to obtain the start port number. Run the `cat /proc/sys/net/ipv4/ip_local_port_range |cut -f 2` command to obtain the end port number. Subtract the start port number from the end port number to obtain the total number of temporary ports. If the total number of temporary ports is less than 28,232, the random port range of the OS is too small. In this case, contact the system administrator to expand the port range.

**Step 5** Run the `ss -ant 2>/dev/null | grep -v LISTEN | awk 'NR > 2 {print $4}'|cut -d ':' -f 2 | awk '$1 > "start port number" {print $1}' | sort -u | wc -l` command to calculate the number of used temporary ports.

**Step 6** Calculate the usage of temporary ports using the following formula: Usage of temporary ports = (Number of used temporary ports/Total number of temporary ports) x 100. Check whether the usage exceeds the threshold.

- If yes, go to [Step 8](#).
- If no, go to [Step 7](#).

**Step 7** Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 8](#).

**Check whether the system environment is normal.**

**Step 8** Run the following command to import the temporary file and view the frequently used ports in the `port_result.txt` file:

```
netstat -tnp > $BIGDATA_HOME/tmp/port_result.txt
```

```
netstat -tnp
Active Internet connections (w/o servers)

Proto Recv Send LocalAddress ForeignAddress State PID/ProgramName tcp 0 0 10-120-85-154:45433
10-120-8:25009 CLOSE_WAIT 94237/java
tcp 0 0 10-120-85-154:45434 10-120-8:25009 CLOSE_WAIT 94237/java
tcp 0 0 10-120-85-154:45435 10-120-8:25009 CLOSE_WAIT 94237/java
...
```

**Step 9** Run the following command to check the processes that occupy a large number of ports:

```
ps -ef |grep PID
```

#### NOTE

- `PID` indicates the process ID of the port queried in [Step 8](#).
- Run the following command to collect information about all processes in the system and check the processes that occupy a large number of ports:

```
ps -ef > $BIGDATA_HOME/tmp/ps_result.txt
```

**Step 10** Contact the system administrator to clear the processes that occupy a large number of ports. Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Step 11** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.446 ALM-12053 File Handle Usage Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the handler usage every 30 seconds. This alarm is generated when the handle usage exceeds the threshold (the default threshold is **80%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System > Threshold Configuration > Device > Host > Host Status > Host File Handle Usage > Host File Handle Usage**.

If the **hit number** is **1**, this alarm is cleared when the file handle usage is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the file handle usage is less than or equal to 90% of the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12053    | Major          | Yes        |

### Parameters

| Parameter         | Description                                             |
|-------------------|---------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated. |
| RoleName          | Specifies the role for which the alarm is generated.    |
| HostName          | Specifies the host for which the alarm is generated.    |
| Trigger Condition | Specifies the threshold for triggering the alarm.       |

### Impact on the System

The system applications fail to open files, access networks, and perform other I/O operations. The applications are running improperly.



## Possible Causes

- The number of file handles does not meet service requirements.
- The system is abnormal.

## Procedure

### Increase the number of file handles.

- Step 1** Go to the MRS cluster details page and choose **Alarms**.
- Step 2** In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address of the host for which the alarm is generated.
- Step 3** Use PuTTY to log in to the host for which the alarm is generated as user **root**.
- Step 4** Run the **ulimit -n** command to check the maximum number of handles set in the system.
- Step 5** If the file handle usage exceeds the threshold, contact the system administrator to increase the number of system file handles.
- Step 6** Wait 5 minutes and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 7](#).

### Check whether the system environment is normal.

- Step 7** Contact the system administrator to check whether the OS is abnormal.
- If yes, rectify the operating system fault and go to [Step 8](#).
  - If no, go to [Step 9](#).
- Step 8** Wait 5 minutes and check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 9](#).
- Step 9** Collect fault information.
1. On MRS Manager, choose **System > Export Log**.
  2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.447 ALM-12054 Invalid Certificate File (For MRS 2.x or Earlier)

### Description

The system checks whether the certificate file is invalid (has expired or is not yet valid) on 23:00 every day. This alarm is generated when the certificate file is invalid.

This alarm is cleared when the status of the newly imported certificate is valid.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12054    | Major          | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

The system reminds users that the certificate file is invalid. If the certificate file is invalid, some functions are restricted and cannot be used properly.

## Possible Causes

No certificate (HA root certificate or HA user certificate) is imported to the system, the certificate fails to be imported, or the certificate file is invalid.

## Procedure

**Check the alarm cause.**

**Step 1** Go to the MRS cluster details page and choose **Alarms**.

**Step 2** In the real-time alarm list, click the row that contains the alarm.

In the **Alarm Details** area, view the additional information about the alarm.

- If **CA Certificate** is displayed in the additional alarm information, use PuTTY to log in to the active OMS management node as user **omm** and go to [Step 3](#).
- If **HA root Certificate** is displayed in the additional information, check **Location** to obtain the name of the host involved in this alarm. Then use PuTTY to log in to the host as user **omm** and go to [Step 4](#).
- If **HA server Certificate** is displayed in the additional information, check **Location** to obtain the name of the host involved in this alarm. Then use PuTTY to log in to the host as user **omm** and go to [Step 5](#).

**Check the validity period of the certificate files in the system.**

**Step 3** Check whether the current system time is in the validity period of the CA certificate.

Run the **openssl x509 -noout -text -in \${CONTROLLER\_HOME}/security/cert/root/ca.crt** command to check the effective time and due time of the root certificate.

- If yes, go to [Step 8](#).
- If no, go to [Step 6](#).

**Step 4** Check whether the current system time is in the validity period of the HA root certificate.

Run the **openssl x509 -noout -text -in \${CONTROLLER\_HOME}/security/certHA/root-ca.crt** command to check the effective time and due time of the HA root certificate.

- If yes, go to [Step 8](#).
- If no, go to [Step 7](#).

**Step 5** Check whether the current system time is in the validity period of the HA user certificate.

Run the **openssl x509 -noout -text -in \${CONTROLLER\_HOME}/security/certHA/server.crt** command to check the effective time and due time of the HA user certificate.

- If yes, go to [Step 8](#).
- If no, go to [Step 7](#).

The following is an example of the effective time and expiration time of a CA or HA certificate:

```
Certificate:
Data:
 Version: 3 (0x2)
 Serial Number:
 97:d5:0e:84:af:ec:34:d8
 Signature Algorithm: sha256WithRSAEncryption
 Issuer: C=CountryName, ST=State, L=Locality, O=Organization, OU=IT, CN=HADOOP.COM
 Validity
 Not Before: Dec 13 06:38:26 2016 GMT // Effective time
 Not After : Dec 11 06:38:26 2026 GMT // Expiration time
```

### Import certificate files.

**Step 6** Import a new CA certificate file.

Contact O&M personnel to apply for or generate a new CA certificate file and import it. Manually clear the alarm and check whether this alarm is generated again during periodic check.

#### NOTE

If the Ranger certificate has expired, see [How Do I Update the Ranger Certificate?](#) to rectify the fault.

- If yes, go to [Step 8](#).
- If no, no further action is required.

**Step 7** Import a new HA certificate file.

Apply for or generate a new HA certificate file and import it by referring to [Replacing an HA Certificate](#). Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to [Step 8](#).
- If no, no further action is required.

**Step 8** Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

For details about how to handle an expired OBS certificate, see [Expired OBS Certificate in a Cluster](#).

## 7.12.448 ALM-12055 Certificate File Is About to Expire (For MRS 2.x or Earlier)

### Description

The system checks the certificate file on 23:00 every day. This alarm is generated if the certificate file is about to expire with a validity period less than days set in the alarm threshold.

This alarm is generated if the status of the newly imported certificate is valid.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12055    | Minor          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

The system reminds users that the license is about to expire. If the license expires, some functions are restricted and cannot be used properly.

## Possible Causes

The remaining validity period of the CA certificate, HA root certificate, or HA user certificate is smaller than the alarm threshold.

## Procedure

**Check the alarm cause.**

**Step 1** Go to the MRS cluster details page and choose **Alarms**.

**Step 2** In the real-time alarm list, click the row that contains the alarm.

In the **Alarm Details** area, view the additional information about the alarm.

- If **CA Certificate** is displayed in the additional alarm information, use PuTTY to log in to the active OMS management node as user **omm** and go to [Step 3](#).
- If **HA root Certificate** is displayed in the additional information, check **Location** to obtain the name of the host involved in this alarm. Then use PuTTY to log in to the host as user **omm** and go to [Step 4](#).
- If **HA server Certificate** is displayed in the additional information, check **Location** to obtain the name of the host involved in this alarm. Then use PuTTY to log in to the host as user **omm** and go to [Step 5](#).

**Check the validity period of the certificate files in the system.**

**Step 3** Check whether the remaining validity period of the CA certificate is smaller than the alarm threshold.

Run the `openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/cert/root/ca.crt` command to check the effective time and due time of the root certificate.

- If yes, go to [Step 6](#).
- If no, go to [Step 8](#).

**Step 4** Check whether the remaining validity period of the HA root certificate is smaller than the alarm threshold.

Run the `openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/root-ca.crt` command to check the effective time and due time of the HA root certificate.

- If yes, go to [Step 7](#).
- If no, go to [Step 8](#).

**Step 5** Check whether the remaining validity period of the HA user certificate is smaller than the alarm threshold.

Run the `openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/server.crt` command to check the effective time and due time of the HA user certificate.

- If yes, go to [Step 7](#).
- If no, go to [Step 8](#).

The following is an example of the effective time and expiration time of a CA or HA certificate:

```
Certificate:
Data:
 Version: 3 (0x2)
 Serial Number:
 97:d5:0e:84:af:ec:34:d8
 Signature Algorithm: sha256WithRSAEncryption
 Issuer: C=CountryName, ST=State, L=Locality, O=Organization, OU=IT, CN=HADOOP.COM
 Validity
 Not Before: Dec 13 06:38:26 2016 GMT // Effective time
 Not After : Dec 11 06:38:26 2026 GMT // Expiration time
```

### Import certificate files.

**Step 6** Import a new CA certificate file.

Contact O&M personnel to apply for or generate a new CA certificate file and import it. Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to [Step 8](#).
- If no, no further action is required.

**Step 7** Import a new HA certificate file.

Apply for or generate a new HA certificate file and import it by referring to [Replacing an HA Certificate](#). Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to [Step 8](#).
- If no, no further action is required.

**Step 8** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

For details about how to handle an expired OBS certificate, see [Expired OBS Certificate in a Cluster](#).

## 7.12.449 ALM-12180 Disk Card I/O (For MRS 2.x or Earlier)

### Description

**For MRS 2.x or earlier:**

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - The system collects data every 3 seconds, and detects that the **svctm** value exceeds 6s for 10 consecutive periods within 30 seconds.
  - The system collects data every 3 seconds, and detects that the **avgqu-sz** value is greater than 0, the IOPS or bandwidth is 0, and the **ioutil** value is greater than **99%** for 10 consecutive periods within 30 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - The system collects data every 3 seconds, and detects that the **svctm** value exceeds 2s for 10 consecutive periods within 30 seconds.
  - The system collects data every 3 seconds, and detects that the **avgqu-sz** value is greater than 0, the IOPS or bandwidth is 0, and the **ioutil** value is greater than **99%** for 10 consecutive periods within 30 seconds.

This alarm is automatically cleared when none of the conditions are met for 90 seconds.

#### For MRS 1.9.3.10 or later:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The **svctm** latency reaches 6 seconds within 30 seconds in at least seven collection periods.
  - The system collects data every 3 seconds by default, and within 30 seconds, at least 10 collection cycles have a disk queue depth (**avgqu-sz**) > 0, IOPS = 0, or bandwidth = 0, and **ioutil** > 99%.
  - By default, the system collects data every 3 seconds. At least 50% of detected **svctm** take no less than 1000 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The **svctm** latency reaches 3 seconds within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. Disk queue depth (**avgqu-sz**) > 0 and IOPS = 0, or bandwidth = 0 and **ioutil** > 99% in at least 10 collection periods within 30 seconds.
  - By default, the system collects data every 3 seconds. At least 50% of detected **svctm** take no less than 500 ms within 300 seconds.

The collection period is 3 seconds, and the detection period is 30 or 300 seconds. This alarm is automatically cleared when none of the conditions are met for three consecutive detection periods (30 or 300 seconds).

#### NOTE

For details about how to obtain related parameters, see [Related Information](#).

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12180    | Major          | Yes        |

## Parameters

| Parameter   | Description                                                       |
|-------------|-------------------------------------------------------------------|
| Source      | Specifies the cluster or system for which the alarm is generated. |
| ServiceName | Specifies the service for which the alarm is generated.           |
| RoleName    | Specifies the role for which the alarm is generated.              |
| HostName    | Specifies the host for which the alarm is generated.              |
| DiskName    | Specifies the disk for which the alarm is generated.              |

## Impact on the System

A continuously high I/O usage may adversely affect service operations and result in service loss.

## Possible Causes

The disk is aged.

## Procedure

### Replace the disk.

**Step 1** Log in to FusionInsight Manager and choose **O&M > Alarm > Alarms**.

**Step 2** View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is reported.

**Step 3** Replace the faulty disk.

**Step 4** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, perform [Step 5](#).

### Collect the fault information.

**Step 5** On MRS Manager, choose **System > Export Log**.

**Step 6** Contact O&M engineers and send the collected logs.

----End

## Alarm Clearing

This alarm is automatically cleared after the fault is rectified.



## Related Information

To obtain the related parameters, perform the following steps:

- Run the following command in the OS to collect data:

**iotstat -x -t 1 1**

```

[root@ ~]# iostat -x -t 1 1
Linux 4.18.0-147.5.2.el8.x86_64 (node-master1cxy) 10/12/2022 _x86_64_ (8 CPU)

10/12/2022 05:24:00 PM
avg-cpu: %user %nice %system %iowait %steal %idle
 24.49 0.00 13.82 0.11 0.00 61.58

Device r/s kB/s rrrq/s %rrqm r_await rareq-sz w/s kB/s wrqm/s %wrqm w_await wareq-sz d/s kB/s drqm/s %drqm d_await dareq-sz aqu-sz %util
dm-0 1.59 57.23 0.00 0.00 1.22 35.94 15.80 124.80 0.00 0.00 2.59 7.90 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.04 0.79
vda 0.97 3.28 0.00 0.00 0.67 4.41 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.01
vda 1.90 61.59 0.02 0.96 1.05 32.43 22.16 403.26 33.50 60.19 1.80 18.20 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.03 1.80
vdb 0.11 2.51 0.00 0.01 0.68 22.22 24.05 351.18 16.74 41.03 1.02 14.60 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.01 1.59

```

The command parameters are as follows:

**avgqu-sz** indicates the disk queue depth.

The sum of **r/s** and **w/s** is the IOPS.

The sum of **rkB/s** and **wkB/s** is the bandwidth.

**%util** is the value of **ioutil**.

- The value of **svctm** is calculated as follows:

$$svctm = (tot\_ticks\_new - tot\_ticks\_old) / (rd\_ios\_new + wr\_ios\_new - rd\_ios\_old - wr\_ios\_old)$$

**For MRS 2.x or earlier:**

If **rd\_ios\_new + wr\_ios\_new - rd\_ios\_old - wr\_ios\_old** is 0, then **svctm** is 0.

**For MRS 1.9.3.10 or later:**

When the detection period is 30 seconds, if **rd\_ios\_new + wr\_ios\_new - rd\_ios\_old - wr\_ios\_old = 0**, then **svctm = 0**.

When the detection period is 300 seconds and **rd\_ios\_new + wr\_ios\_new - rd\_ios\_old - wr\_ios\_old = 0**, if **tot\_ticks\_new - tot\_ticks\_old = 0**, then **svctm = 0**; otherwise, the value of **svctm** is infinite.

The parameters in the preceding expression can be obtained as follows:

Obtain the parameter values from the data collected via the **cat /proc/diskstats** command run by the system every 3 seconds. The following shows an example.

```

[comm@ ~]# cat /proc/diskstats
253 0 vda 1101553 35446 83439787 3338546 28744856 48314024 1054257652 52667332 0 19569526 10342913 0 0 0 0
253 1 vda1 590970 25494 54533791 2565698 3446015 6750402 215791076 121115169 0 6474429 11339985 0 0 0 0
253 2 vda2 15 0 108 136 0 0 0 0 150 129 0 0 0 0
253 5 vda5 22373 1364 2525122 79502 4212374 4104759 161597984 8145606 0 3598808 6239095 0 0 0 0
253 6 vda6 11145 314 529002 85050 259201 70368 4412408 321454 0 189336 259725 0 0 0 0
253 7 vda7 157987 105 3477434 149542 6507077 1028968 140666992 14349866 0 1679035 11116587 0 0 0 0
253 8 vda8 312935 8169 22369722 458354 12179958 34360589 531802640 17724858 0 9060731 11385470 0 0 0 0
253 16 vdb 275920 21939 15977738 2171665 39472291 28236575 2653825040 482230505 0 30580346 465962048 0 0 0 0
253 17 vdb1 275439 21939 15948866 2171472 31290400 28236555 2653824832 481837775 0 30036724 465855080 0 0 0 0
7 0 loop0 356 0 17442 150 0 0 0 0 149 105 0 0 0 0
[comm@ ~]# cat /proc/diskstats
253 0 vda 1101553 35446 83439787 3338546 28747977 48319338 1054352084 52672715 0 19571460 40346640 0 0 0 0
253 1 vda1 590970 25494 54533791 2565698 3446015 6750402 215791076 12115169 0 6474429 11339985 0 0 0 0
253 2 vda2 15 0 108 136 0 0 0 0 150 129 0 0 0 0
253 5 vda5 22373 1364 2525122 79502 4212822 4105244 161614088 8146153 0 3599216 6239432 0 0 0 0
253 6 vda6 11145 314 529002 85050 259245 70433 4413368 321489 0 189389 259730 0 0 0 0
253 7 vda7 157987 105 3477434 149542 6507759 1029060 140677872 14351373 0 1679157 11117724 0 0 0 0
253 8 vda8 312935 8169 22369722 458354 12181277 34364199 531855680 1772525 0 9061647 11387424 0 0 0 0
253 16 vdb 275920 21939 15977738 2171665 39477604 28238831 2653881640 482234435 0 30581946 465964144 0 0 0 0
253 17 vdb1 275439 21939 15948866 2171472 31293358 28238811 2653881432 481841639 0 30038274 465857164 0 0 0 0
7 0 loop0 356 0 17442 150 0 0 0 0 149 105 0 0 0 0

```

In the data collected for the first time, the number in the fourth column is the value of **rd\_ios\_old**, the number in the eighth column is the value of **wr\_ios\_old**, and the number in the thirteenth column is the value of **tot\_ticks\_old**.

In the data collected for the second time, the number in the fourth column is the value of **rd\_ios\_new**, the number in the eighth column is the value of

**wr\_ios\_new**, and the number in the thirteenth column is the value of **tot\_ticks\_new**.

In this case, the value of **svctm** is as follows:

$$(19571460 - 19569526) / (1101553 + 28747977 - 1101553 - 28744856) = 0.6197$$

## 7.12.450 ALM-12357 Failed to Export Audit Logs to OBS (For MRS 2.x or Earlier)

### Description

If the user has configured audit log export to the OBS on MRS Manager, the system regularly exports audit logs to the OBS. This alarm is reported if the system fails to access the OBS.

This alarm is cleared after the system exports audit logs to the OBS successfully.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 12357    | Major          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

### Impact on the System

The local system saves a maximum of seven compressed service audit log files. If this alarm persists, local service audit logs may be lost.

The local system saves a maximum of 50 management audit log files (each file contains 100,000 records). If this alarm persists, local management audit logs may be lost.

### Possible Causes

- Connection to the OBS server fails.
- The specified OBS file system does not exist.

- The user AK/SK information is invalid.
- The local OBS configuration cannot be obtained.

## Procedure

- Step 1** Log in to the OBS server and check whether the OBS server can be properly accessed.
- If yes, go to [Step 3](#).
  - If no, go to [Step 2](#).
- Step 2** Contact the maintenance personnel to repair OBS. Then check whether the alarm is cleared.
- If yes, no further action is required.
  - If no, go to [Step 3](#).
- Step 3** On MRS Manager, choose **System > Export Audit Log**. Check whether the AK/SK information, file system name, and path are correct.
- If yes, go to [Step 5](#).
  - If no, go to [Step 4](#).
- Step 4** Correct the information. Then check whether the alarm is cleared when the export task is executed again.

### NOTE

To check alarm clearance quickly, you can set the start time of audit log collection to 10 or 30 minutes later than the current time. After checking the result, restore the original start time.

- If yes, no further action is required.
- If no, go to [Step 5](#).

- Step 5** Collect fault information.
1. On MRS Manager, choose **System > Export Log**.
  2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.451 ALM-13000 ZooKeeper Service Unavailable (For MRS 2.x or Earlier)

### Description

The system checks the ZooKeeper service status every 30 seconds. This alarm is generated when the ZooKeeper service is unavailable.

This alarm is cleared when the ZooKeeper service recovers.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 13000    | Critical       | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

ZooKeeper fails to provide coordination services for upper-layer components and the components depending on ZooKeeper may not run properly.

## Possible Causes

- The ZooKeeper instance is abnormal.
- The disk capacity is insufficient.
- The network is faulty.
- The DNS is installed on the ZooKeeper node.

## Procedure

### Check the ZooKeeper service instance status.

- Step 1** On the MRS cluster details page, choose **Components > ZooKeeper > quorumpeer**.
- Step 2** Check whether the ZooKeeper instances are normal.
  - If yes, go to [Step 6](#).
  - If no, go to [Step 3](#).
- Step 3** Select instances whose status is not good and choose **More > Restart Instance**.
- Step 4** Check whether the instance status is good after restart.
  - If yes, go to [Step 5](#).
  - If no, go to [Step 19](#).
- Step 5** On the **Alarms** tab page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6](#).

**Check disk status.**

**Step 6** On the MRS cluster details page, choose **Components > ZooKeeper > quorumpeer**, and check the host information of each node housing the ZooKeeper instance.

**Step 7** On the MRS cluster details page, click the **Nodes** tab and expand a node group.

**Step 8** In the **Disk Usage** column, check whether the disk space of each node housing ZooKeeper instances is insufficient (disk usage exceeds 80%).

- If yes, go to [Step 9](#).
- If no, go to [Step 11](#).

**Step 9** Expand the disk capacity. For details, see [ALM-12017 Insufficient Disk Capacity \(For MRS 2.x or Earlier\)](#).

**Step 10** On the **Alarms** tab page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 11](#).

**Check network communication status.**

**Step 11** On the Linux node housing the ZooKeeper instance, run the **ping** command to check whether the host names of other nodes housing the ZooKeeper instances can be pinged successfully.

- If yes, go to [Step 15](#).
- If no, go to [Step 12](#).

**Step 12** Modify the IP addresses in **/etc/hosts** and add the mapping between host names and IP addresses.

**Step 13** Run the **ping** command again to check whether the host names of other nodes housing the ZooKeeper instances can be pinged successfully.

- If yes, go to [Step 14](#).
- If no, go to [Step 19](#).

**Step 14** On the **Alarms** tab page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 15](#).

**Check the DNS.**

**Step 15** Check whether the DNS is installed on the node housing the ZooKeeper instance. On the Linux node housing the ZooKeeper instance, run the **cat /etc/resolv.conf** command to check whether the file is empty.

- If yes, go to [Step 16](#).
- If no, go to [Step 19](#).

**Step 16** Run the **service named status** command to check whether the DNS is started.

- If yes, go to [Step 17](#).

- If no, go to [Step 19](#).

**Step 17** Run the **service named stop** command to stop the DNS service. If "Shutting down name server BIND waiting for named to shut down (28s)" is displayed, the DNS service is stopped successfully. Comment out the content (if any) in **/etc/resolv.conf**.

**Step 18** On the **Alarms** tab page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 19](#).

**Step 19** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.452 ALM-13001 Available ZooKeeper Connections Are Insufficient (For MRS 2.x or Earlier)

### Description

The system checks ZooKeeper connections every 30 seconds. This alarm is generated when the system detects that the number of used ZooKeeper instance connections exceeds the threshold (80% of the maximum connections).

This alarm is cleared when the number of used ZooKeeper instance connections is less than the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 13001    | Major          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| Trigger Condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

Available ZooKeeper connections are insufficient. When the connection usage reaches 100%, external connections cannot be handled.

## Possible Causes

The number of connections to the ZooKeeper node exceeds the threshold. Connection leakage occurs on some connection processes, or the maximum number of connections does not meet the requirement of the actual scenario.

## Procedure

### Step 1 Check the connection status.

1. On the MRS cluster details page, choose **Alarms > ALM-13001 Available ZooKeeper Connections Are Insufficient > Location**. Check the IP address of the node for which the alarm is generated.
2. Obtain the PID of the ZooKeeper process. Log in to the node for which this alarm is generated and run the **pgrep -f proc\_zookeeper** command.
3. Check whether the PID can be successfully obtained.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 2](#).
4. Obtain all the IP addresses connected to the ZooKeeper instance and the number of connections and check 10 IP addresses with top connections. Run the **lsof -i|grep \$pid | awk '{print \$9}' | cut -d : -f 2 | cut -d \> -f 2 | awk '{a[\$1]++} END {for(i in a){print i,a[i] | "sort -r -g -k 2"}}' | head -10** command based on the obtained PID value. (**\$pid** is the PID obtained in the preceding step.)
5. Check whether the node IP addresses and the number of connections are successfully obtained.
  - If yes, go to [Step 1.6](#).
  - If no, go to [Step 2](#).
6. Obtain the ID of the port connected to the process. Run the **lsof -i|grep \$pid | awk '{print \$9}'|cut -d \> -f 2 |grep \$IP| cut -d : -f 2** command based on the obtained PID and IP address. (**\$pid** and **\$IP** are the PID and IP address obtained in the preceding step.)
7. Check whether the port ID is successfully obtained.
  - If yes, go to [Step 1.8](#).
  - If no, go to [Step 2](#).

8. Obtain the ID of the connected process. Log in to each IP address and run the following command based on the obtained port ID: **lsof -i|grep \$port**. (**\$port** is the port ID obtained in the preceding step.)
9. Check whether the process ID is successfully obtained.
  - If yes, go to [Step 1.10](#).
  - If no, go to [Step 2](#).
10. Check whether connection leakage occurs on the process based on the obtained process ID.
  - If yes, go to [Step 1.11](#).
  - If no, go to [Step 1.12](#).
11. Close the process where connection leakage occurs and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 1.12](#).
12. On the MRS cluster details page, choose **Components > ZooKeeper > Service Configuration**. Set **Type** to **All**, choose **quorumpeer > Performance**, and change the value of **maxCnxns** to **20000** or more.
13. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.453 ALM-13002 ZooKeeper Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the ZooKeeper service status every 30 seconds. The alarm is generated when the memory usage of a ZooKeeper instance exceeds the threshold (80% of the maximum memory).

The alarm is cleared when the memory usage is less than the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 13002    | Major          | Yes        |



## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger Condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

If the available ZooKeeper memory is insufficient, a memory overflow occurs and the service breaks down.

## Possible Causes

The memory usage of the ZooKeeper instance is overused or the memory is inappropriately allocated.

## Procedure

### Step 1 Check the memory usage.

1. On the MRS cluster details page, choose **Alarms > ALM-13002 ZooKeeper Memory Usage Exceeds the Threshold > Location**. Check the IP address of the instance for which the alarm is generated.
2. On the MRS cluster details page, choose **Components > ZooKeeper > Instances > quorumpeer** (IP address of the instance for which the alarm is generated) **> Customize > ZooKeeper Heap And Direct Buffer Resource**. Check the heap memory usage.
3. Check whether the used heap memory of ZooKeeper reaches 80% of the maximum heap memory specified for ZooKeeper.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 1.6](#).
4. On MRS Manager, choose **Services > ZooKeeper > Configuration > All > quorumpeer > System**. Increase the value of **-Xmx** in **GC\_OPTS** as required.
5. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 1.6](#).
6. On the MRS cluster details page, choose **Components > ZooKeeper > Instances > quorumpeer** (IP address of the instance for which the alarm is

generated) > **Customize** > **ZooKeeper Heap And Direct Buffer Resource**.  
Check the direct buffer memory usage.

7. Check whether the used direct buffer memory of ZooKeeper reaches 80% of the maximum direct buffer memory specified for ZooKeeper.
  - If yes, go to [Step 1.8](#).
  - If no, go to [Step 2](#).
8. On the MRS cluster details page, choose **Components** > **ZooKeeper** > **Service Configuration**. Set **Type** to **All** and choose **quorumpeer** > **System**.  
Increase the value of **-XX:MaxDirectMemorySize** in **GC\_OPTS** as required.
9. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.454 ALM-14000 HDFS Service Unavailable (For MRS 2.x or Earlier)

### Description

The system checks the service status of NameService every 30 seconds. This alarm is generated when the system considers that the HDFS service is unavailable because all the NameService services are abnormal.

This alarm is cleared when at least one NameService service is normal and the system considers that the HDFS service recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14000    | Critical       | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |

| Parameter | Description                                          |
|-----------|------------------------------------------------------|
| RoleName  | Specifies the role for which the alarm is generated. |
| HostName  | Specifies the host for which the alarm is generated. |

## Impact on the System

HDFS fails to provide services for HDFS service-based upper-layer components, such as HBase and MapReduce. As a result, users cannot read or write files.

## Possible Causes

- ZooKeeper is abnormal.
- All NameService services are abnormal.

## Procedure

### Step 1 Check the ZooKeeper status.

1. Go to the MRS cluster details page. On the **Components** tab page, check whether the health status of the ZooKeeper service is **Good**.
  - If yes, go to [Step 1.2](#).
  - If no, go to [Step 2.1](#).
2. Rectify the health status of the ZooKeeper service. For details, see [ALM-13000 ZooKeeper Service Unavailable \(For MRS 2.x or Earlier\)](#). Then check whether the health status of the ZooKeeper service is **Good**.
  - If yes, go to [Step 1.3](#).
  - If no, go to [Step 3](#).
3. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

### Step 2 Handle the NameService service exception alarm.

1. Go to the MRS cluster details page. On the **Alarms** page, check whether all NameService services have abnormal alarms.
  - If yes, go to [Step 2.2](#).
  - If no, go to [Step 3](#).
2. Handle the abnormal NameService services following the instructions in [ALM-14010 NameService Is Abnormal \(For MRS 2.x or Earlier\)](#) and check whether each NameService service exception alarm is cleared.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 3](#).
3. Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 3](#).

**Step 3** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

**Reference**

None

## 7.12.455 ALM-14001 HDFS Disk Usage Exceeds the Threshold (For MRS 2.x or Earlier)

**Description**

The system checks the disk usage of the HDFS cluster every 30 seconds and compares the actual disk usage with the threshold. The HDFS cluster disk usage indicator has a default threshold. This alarm is generated when the HDFS disk usage exceeds the threshold.

This alarm is cleared when the disk usage of the HDFS cluster is less than or equal to the threshold.

**Attribute**

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14001    | Major          | Yes        |

**Parameters**

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| NSName            | Specifies the NameService service for which the alarm is generated.                 |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

The performance of writing data to HDFS is affected.

## Possible Causes

The disk space configured for the HDFS cluster is insufficient.

## Procedure

**Step 1** Check the disk capacity and delete unnecessary files.

1. On the MRS cluster details page, choose **Components** > **HDFS**. The **Service Status** page is displayed.
2. In the **Charts** area, view the value of the monitoring indicator **Percentage of HDFS Capacity** to check whether the HDFS disk usage exceeds the threshold (80% by default).
  - If yes, go to **Step 1.3**.
  - If no, go to **Step 3**.
3. Use the client on the cluster node and run the **hdfs dfsadmin -report** command to check whether the value of **DFS Used%** is less than 100% minus the threshold.
  - If yes, go to **Step 1.5**.
  - If no, go to **Step 3**.
4. Use the client on the cluster node and run the **hdfs dfs -rm -r file or directory path** command to delete unnecessary files.
5. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 2.1**.

**Step 2** Expand the system.

1. Expand the disk capacity.
2. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 3**.

**Step 3** Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.456 ALM-14002 DataNode Disk Usage Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the DataNode disk usage every 30 seconds and compares the actual disk usage with the threshold. The **Percentage of DataNode Capacity** indicator has a default threshold. This alarm is generated when the value of the **Percentage of DataNode Capacity** indicator exceeds the threshold.

This alarm is cleared when the value of the **Percentage of DataNode Capacity** indicator is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14002    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

### Impact on the System

Insufficient disk space will impact read/write to HDFS.

### Possible Causes

- The disk space configured for the HDFS cluster is insufficient.
- Data skew occurs among DataNodes.

### Procedure

**Step 1** Check the cluster disk capacity.

1. Go to the MRS cluster details page. On the **Alarms** page, check whether the ALM-14001 HDFS Disk Usage Exceeds the Threshold alarm exists.

- If yes, go to [Step 1.2](#).
  - If no, go to [Step 2.1](#).
2. Handle the alarm by following the instructions in ALM-14001 HDFS Disk Usage Exceeds the Threshold and check whether the alarm is cleared.
    - If yes, go to [Step 1.3](#).
    - If no, go to [Step 3](#).
  3. Wait 5 minutes and check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to [Step 2.1](#).

**Step 2** Check the balance status of DataNodes.

1. Use the client on the cluster node, run the **hdfs dfsadmin -report** command to view the value of **DFS Used%** on the DataNode for which the alarm is generated, and compare the value with those on other DataNodes. Check whether the difference between the values is larger than 10.
  - If yes, go to [Step 2.2](#).
  - If no, go to [Step 3](#).
2. If data skew occurs, use the client on the cluster node and run the **hdfs balancer -threshold 10** command.
3. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3](#).

**Step 3** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.457 ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the number of lost blocks every 30 seconds and compares the number of lost blocks with the threshold. The lost blocks indicator has a default threshold. This alarm is generated when the number of lost blocks exceeds the threshold.

This alarm is cleared when the number of lost blocks is less than or equal to the threshold.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14003    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| NSName            | Specifies the NameService service for which the alarm is generated.                 |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

Data stored in HDFS is lost. HDFS may enter the safe mode and cannot provide write services. Lost block data cannot be restored.

## Possible Causes

- The DataNode instance is abnormal.
- Data is deleted.

## Procedure

**Step 1** Check the DataNode instance.

1. On the MRS cluster details page, choose **Components > HDFS > Instances**.
2. Check whether the status of all DataNode instances is **Good**.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 1.3](#).
3. Restart the DataNode instance and check whether the restart is successful.
  - If yes, go to [Step 2.2](#).
  - If no, go to [Step 2.1](#).

**Step 2** Delete the damaged file.



1. Use the client on the cluster node. Run the **hdfs fsck / -delete** command to delete the lost file. Then rewrite the file and recover the data.
2. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3](#).

**Step 3** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.458 ALM-14004 Number of Damaged HDFS Blocks Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the number of damaged blocks every 30 seconds and compares the number of damaged blocks with the threshold. The damaged blocks indicator has a default threshold. This alarm is generated when the number of damaged blocks exceeds the threshold.

This alarm is cleared when the number of damaged blocks is less than or equal to the threshold. You are advised to run the **hdfs fsck /** command to check whether any file is completely damaged.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14004    | Major          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| NSName            | Specifies the NameService service for which the alarm is generated.                 |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

Data is damaged and HDFS fails to read files.

## Possible Causes

- The DataNode instance is abnormal.
- Data verification information is damaged.

## Procedure

**Step 1** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.459 ALM-14006 Number of HDFS Files Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system periodically checks the number of HDFS files every 30 seconds and compares the number of HDFS files with the threshold. This alarm is generated when the system detects that the number of HDFS files exceeds the threshold.

This alarm is cleared when the number of HDFS files is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14006    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| NSName            | Specifies the NameService service for which the alarm is generated.                 |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

Disk storage space is insufficient, which may result in data import failure. The performance of the HDFS system is affected.

## Possible Causes

The number of HDFS files exceeds the threshold.

## Procedure

**Step 1** Check whether unnecessary files exist in the system.

1. Use the client on the cluster node and run the **hdfs dfs -ls file or directory path** command to check whether the file or directory can be deleted.
  - If yes, go to [Step 1.2](#).
  - If no, go to [Step 2.1](#).
2. Run the **hdfs dfs -rm -r file or directory path** command. Delete unnecessary files, wait 5 minutes, and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

**Step 2** Check the number of files in the system.

1. On MRS Manager, choose **System > Threshold Configuration**.
2. In the navigation tree on the left, choose **Services > HDFS > HDFS File > Total Number of Files**.
3. In the right pane, modify the threshold in the rule based on the number of current HDFS files.  
To check the number of HDFS files, choose **Services > HDFS**, click **Customize** in the **Real-Time Statistics** area on the right, and select the **HDFS File** monitoring item.

4. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3](#).

**Step 3** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.460 ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the HDFS NameNode memory usage every 30 seconds and compares the actual memory usage with the threshold. The HDFS NameNode memory usage has a default threshold. This alarm is generated when the HDFS NameNode memory usage exceeds the threshold.

This alarm is cleared when the HDFS NameNode memory usage is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14007    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

If the memory usage of the HDFS NameNode is too high, data read/write performance of HDFS will be affected.

## Possible Causes

The HDFS NameNode memory is insufficient.

## Procedure

**Step 1** Delete unnecessary files.

1. Use the client on the cluster node and run the `hdfs dfs -rm -r file or directory path` command to delete unnecessary files.
2. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.461 ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the HDFS DataNode memory usage every 30 seconds and compares the actual memory usage with the threshold. The HDFS DataNode memory usage has a default threshold. This alarm is generated when the HDFS DataNode memory usage exceeds the threshold.

This alarm is cleared when the HDFS DataNode memory usage is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14007    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

The HDFS DataNode memory usage is too high, which affects the data read/write performance of the HDFS.

## Possible Causes

The HDFS DataNode memory is insufficient.

## Procedure

**Step 1** Delete unnecessary files.

1. Use the client on the cluster node and run the **hdfs dfs -rm -r file or directory path** command to delete unnecessary files.
2. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.462 ALM-14009 Number of Faulty DataNodes Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system periodically checks the number of faulty DataNodes in the HDFS cluster every 30 seconds, and compares the number with the threshold. The number of faulty DataNodes has a default threshold. This alarm is generated when the number of faulty DataNodes in the HDFS cluster exceeds the threshold.

This alarm is cleared when the number of faulty DataNodes in the HDFS cluster is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14009    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

### Impact on the System

Faulty DataNodes cannot provide HDFS services.

### Possible Causes

- DataNodes are faulty or overloaded.
- The network between the NameNode and the DataNode is disconnected or busy.
- NameNodes are overloaded.

## Procedure

### Step 1 Check whether DataNodes are faulty.

1. Use the client on the cluster node to run the **hdfs dfsadmin -report** command to check whether DataNodes are faulty.
  - If yes, go to [Step 1.2](#).
  - If no, go to [Step 2.1](#).
2. On the MRS cluster details page, choose **Components > HDFS > Instances** to check whether the DataNode is stopped.
  - If yes, go to [Step 1.3](#).
  - If no, go to [Step 2.1](#).
3. Select the DataNode instance, and choose **More > Restart Instance** to restart it. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

### Step 2 Check the status of the network between the NameNode and the DataNode.

1. Log in to the service IP address of the node where the faulty DataNode is located, and run the **ping IP address of the NameNode** command to check whether the network between the DataNode and the NameNode is abnormal.
  - If yes, go to [Step 2.2](#).
  - If no, go to [Step 3.1](#).
2. Rectify the network fault. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3.1](#).

### Step 3 Check whether the DataNode is overloaded.

1. On the MRS cluster details page, click **Alarms** and check whether the alarm ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold exists.
  - If yes, go to [Step 3.2](#).
  - If no, go to [Step 4.1](#).
2. Follow procedures in [ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold \(For MRS 2.x or Earlier\)](#) to handle the alarm and check whether the alarm is cleared.
  - If yes, go to [Step 3.3](#).
  - If no, go to [Step 4.1](#).
3. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4.1](#).

### Step 4 Check whether the NameNode is overloaded.

1. On the MRS cluster details page, click **Alarms** and check whether the alarm ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold exists.



- If yes, go to [Step 4.2](#).
  - If no, go to [Step 5](#).
2. Follow procedures in [ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold \(For MRS 2.x or Earlier\)](#) to handle the alarm and check whether the alarm is cleared.
    - If yes, go to [Step 4.3](#).
    - If no, go to [Step 5](#).
  3. Wait 5 minutes and check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to [Step 5](#).

**Step 5** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.463 ALM-14010 NameService Is Abnormal (For MRS 2.x or Earlier)

### Description

The system checks the NameService service status every 180 seconds. This alarm is generated when the NameService service is unavailable.

This alarm is cleared when the NameService service recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14010    | Major          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |

| Parameter | Description                                                         |
|-----------|---------------------------------------------------------------------|
| HostName  | Specifies the host for which the alarm is generated.                |
| NSName    | Specifies the NameService service for which the alarm is generated. |

## Impact on the System

HDFS fails to provide services for upper-layer components based on the NameService service, such as HBase and MapReduce. As a result, users cannot read or write files.

## Possible Causes

- The JournalNode is faulty.
- The DataNode is faulty.
- The disk capacity is insufficient.
- The NameNode enters safe mode.

## Procedure

**Step 1** Check the status of the JournalNode instance.

1. On the MRS Manager home page, click **Components**.
2. Click **HDFS**.
3. Click **Instance**.
4. Check whether the **Health Status** of the JournalNode is **Good**.
  - If yes, go to [Step 2.1](#).
  - If no, go to [Step 1.5](#).
5. Select the faulty JournalNode, and choose **More > Restart Instance**. Check whether the JournalNode successfully restarts.
  - If yes, go to [Step 1.6](#).
  - If no, go to [Step 5](#).
6. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

**Step 2** Check the status of the DataNode instance.

1. On the MRS cluster details page, click **Components**.
2. Click **HDFS**.
3. In **Operation and Health Summary**, check whether the **Health Status** of all DataNodes is **Good**.
  - If yes, go to [Step 3.1](#).
  - If no, go to [Step 2.4](#).

4. Click **Instances**. On the DataNode management page, select the faulty DataNode, and choose **More > Restart Instance**. Check whether the DataNode successfully restarts.
  - If yes, go to [Step 2.5](#).
  - If no, go to [Step 3.1](#).
5. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4.1](#).

**Step 3** Check the disk status.

1. On the MRS cluster details page, click the **Nodes** tab and expand a node group.
2. In the **Disk Usage** column, check whether disk space is insufficient.
  - If yes, go to [Step 3.3](#).
  - If no, go to [Step 4.1](#).
3. Expand the disk capacity.
4. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4.1](#).

**Step 4** Check whether NameNode is in the safe mode.

1. Use the client on the cluster node, and run the **hdfs dfsadmin -safemode get** command to check whether **Safe mode is ON** is displayed.  
Information behind **Safe mode is ON** is alarm information and is displayed based actual conditions.
  - If yes, go to [Step 4.2](#).
  - If no, go to [Step 5](#).
2. Use the client on the cluster node and run the **hdfs dfsadmin -safemode leave** command.
3. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 5](#).

**Step 5** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.464 ALM-14011 HDFS DataNode Data Directory Is Not Configured Properly (For MRS 2.x or Earlier)

### Description

The DataNode parameter **dfs.datanode.data.dir** specifies the DataNode data directory. This alarm is generated in any of the following scenarios:

- A configured data directory cannot be created.
- A data directory uses the same disk as other critical directories in the system.
- Multiple directories use the same disk.

This alarm is cleared when the DataNode data directory is configured properly and this DataNode is restarted.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14011    | Major          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

### Impact on the System

If the DataNode data directory is mounted on critical directories such as the root directory, the disk space of the root directory will be used up after running for a long time. This causes a system fault.

If the DataNode data directory is not configured properly, HDFS performance will deteriorate.

### Possible Causes

- The DataNode data directory fails to be created.
- The DataNode data directory uses the same disk as critical directories, such as / or **/boot**.
- Multiple directories in the DataNode data directory use the same disk.

## Procedure

- Step 1** Check the alarm cause and information about the DataNode for which the alarm is generated.
1. On the MRS cluster details page, click **Alarms**. In the alarm list, click the alarm.
  2. In the **Alarm Details** area, view **Alarm Cause** to obtain the cause of the alarm. In **HostName** of **Location**, obtain the host name of the DataNode for which the alarm is generated.
- Step 2** Delete directories that do not comply with the disk plan from the DataNode data directory.
1. Choose **Components > HDFS > Instances**. In the instance list, click the DataNode instance on the node for which the alarm is generated.
  2. Click **Instance Configuration** and view the value of the DataNode parameter **dfs.datanode.data.dir**.
  3. Check whether all DataNode data directories are consistent with the disk plan.
    - If yes, go to [Step 2.4](#).
    - If no, go to [Step 2.7](#).
  4. Modify the DataNode parameter **dfs.datanode.data.dir** and delete the incorrect directories.
  5. Choose **Components > HDFS > Instances** to restart the DataNode instance.
  6. Check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to [Step 2.7](#).
  7. Log in to the DataNode for which the alarm is generated.
    - If the alarm cause is "The DataNode data directory fails to be created", go to [Step 3.1](#).
    - If the alarm cause is "The DataNode data directory uses the same disk as critical directories, such / or /boot", go to [Step 4.1](#).
    - If the alarm cause is "Multiple directories in the DataNode data directory use the same disk", go to [Step 5.1](#).
- Step 3** Check whether the DataNode data directory fails to be created.
1. Run the following commands to switch the user:  
**sudo su - root**  
**su - omm**
  2. Run the **ls** command to check whether the directories exist in the DataNode data directory.
    - If yes, go to [Step 7](#).
    - If no, go to [Step 3.3](#).
  3. Run the **mkdir data directory** command to create a directory and check whether the directory is successfully created.
    - If yes, go to [Step 6.1](#).

- If no, go to [Step 3.4](#).
  - 4. Click **Alarms** to check whether alarm ALM-12017 Insufficient Disk Capacity exists.
    - If yes, go to [Step 3.5](#).
    - If no, go to [Step 3.6](#).
  - 5. Adjust the disk capacity and check whether alarm ALM-12017 Insufficient Disk Capacity is cleared. For details, see [ALM-12017 Insufficient Disk Capacity \(For MRS 2.x or Earlier\)](#).
    - If yes, go to [ALM-12017 Insufficient Disk Capacity \(For MRS 2.x or Earlier\)](#).
    - If no, go to [Step 7](#).
  - 6. Check whether user **omm** has the **rwX** or **X** permission of all the upper-layer directories of the directory. (For example, for **/tmp/abc/**, user **omm** has the **X** permission for directory **tmp** and the **rwX** permission for directory **abc**.)
    - If yes, go to [Step 6.1](#).
    - If no, go to [Step 3.7](#).
  - 7. Run the **chmod u+rwX path** or **chmod u+X path** command as the **root** user to add the **rwX** or **X** permission to the paths. Then, go to [Step 3.3](#).
- Step 4** Check whether the DataNode data directory uses the same disk as other critical directories in the system.
1. Run the **df** command to obtain the disk mounting information of each directory in the DataNode data directory.
  2. Check whether the directories mounted to the disk are critical directories, such as **/** or **/boot**.
    - If yes, go to [Step 4.3](#).
    - If no, go to [Step 6.1](#).
  3. Change the value of the DataNode parameter **dfs.datanode.data.dir** and delete the directories that use the same disk as critical directories.
  4. Go to [Step 6.1](#).
- Step 5** Check whether multiple directories in the DataNode data directory use the same disk.
1. Run the **df** command to obtain the disk mounting information of each directory in the DataNode data directory. Record the mounted directory in the command output.
  2. Modify the DataNode node parameter **dfs.datanode.data.dir** to reserve one of the directories mounted on the same disk directory.
  3. Go to [Step 6.1](#).
- Step 6** Restart the DataNode and check whether the alarm is cleared.
1. Choose **Components > HDFS > Instances** to restart the DataNode instance.
  2. Check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to [Step 7](#).
- Step 7** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.465 ALM-14012 HDFS Journalnode Data Is Not Synchronized (For MRS 2.x or Earlier)

### Description

On the active NameNode, the system checks data synchronization on all JournalNodes in the cluster every 5 minutes. This alarm is generated when data on a JournalNode is not synchronized with that on other JournalNodes.

This alarm is cleared in 5 minutes after data on JournalNodes is synchronized.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 14012    | Major          | Yes        |

### Parameters

| Parameter   | Description                                                                                    |
|-------------|------------------------------------------------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated.                                        |
| RoleName    | Specifies the role for which the alarm is generated.                                           |
| IP          | Specifies the service IP address of the JournalNode instance for which the alarm is generated. |

### Impact on the System

When a JournalNode is working incorrectly, data on the node is not synchronized with that on other JournalNodes. If data on more than half of JournalNodes is not synchronized, the NameNode cannot work correctly, making the HDFS service unavailable.

## Possible Causes

- The JournalNode instance has not been started or has been stopped.
- The JournalNode instance is working incorrectly.
- The network of the JournalNode is unreachable.

## Procedure

**Step 1** Check whether the JournalNode instance has been started.

1. On the MRS cluster details page, click **Alarms**. In the alarm list, click the alarm.
2. In the **Alarm Details** area, check **Location** and obtain the IP address of the JournalNode for which the alarm is generated.
3. Choose **Components > HDFS > Instances**. In the instance list, click the JournalNode for which the alarm is generated and check whether **Operating Status** of the node is **Started**.
  - If yes, go to [Step 2.1](#).
  - If no, go to [Step 1.4](#).
4. Select the JournalNode instance and choose **More > Start Instance** to start it.
5. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 2** Check whether the JournalNode instance is working correctly.

1. Check whether **Health Status** of the JournalNode instance is **Good**.
  - If yes, go to [Step 3.1](#).
  - If no, go to [Step 2.2](#).
2. Select the JournalNode instance and choose **More > Restart Instance** to restart it.
3. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 3** Check whether the network of the JournalNode is reachable.

1. On the MRS cluster details page, choose **Components > HDFS > Instances** to check the service IP address of the active NameNode.
2. Log in to the active NameNode.
3. Run the **ping** command to check whether a timeout occurs or the network between the active NameNode and the JournalNode is unreachable.  
**ping** *service IP address of the JournalNode*
  - If yes, go to [Step 3.4](#).
  - If no, go to [Step 4](#).
4. Contact O&M personnel to rectify the network fault. Wait 5 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.



- If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.466 ALM-16000 Percentage of Sessions Connected to the HiveServer to the Maximum Number Allowed Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the percentage of sessions connected to the HiveServer to the maximum number allowed every 30 seconds. This indicator can be viewed on the Hive service monitoring page. This alarm is generated when the the percentage of sessions connected to the HiveServer to the maximum number allowed exceeds the specified threshold (90% by default).

This alarm can be automatically cleared when the percentage is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16000    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

If a connection alarm is generated, too many sessions are connected to the HiveServer and new connections cannot be created.

## Possible Causes

Too many clients are connected to the HiveServer.

## Procedure

**Step 1** Increase the maximum number of connections to Hive.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **Hive > Service Configuration** and switch **Basic** to **All**.
3. Increase the value of the **hive.server.session.control.maxconnections** configuration item. Suppose the value of the configuration item is A, the threshold is B, and sessions connected to the HiveServer is C. Adjust the value of the configuration item according to  $A \times B > C$ . Sessions connected to the HiveServer can be viewed on the Hive monitoring page.
4. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.467 ALM-16001 Hive Warehouse Space Usage Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the Hive warehouse space usage every 30 seconds. The indicator **Percentage of HDFS Space Used by Hive to the Available Space** can be viewed on the Hive service monitoring page. This alarm is generated when the Hive warehouse space usage exceeds the specified threshold (85% by default).

This alarm is cleared when the Hive warehouse space usage is less than or equal to the threshold. You can reduce the warehouse space usage by expanding the warehouse capacity or releasing the used space.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16001    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

The system fails to write data, which causes data loss.

## Possible Causes

- The upper limit of the HDFS capacity available for Hive is too small.
- The system disk space is insufficient.
- Some data nodes break down.

## Procedure

### Step 1 Expand the system configuration.

1. Analyze the cluster HDFS capacity usage and increase the upper limit of the HDFS capacity available for Hive.

Go to the MRS cluster details page, choose **Components > Hive > Service Configuration**, set **Type** to **All**, search for **hive.metastore.warehouse.size.percent**, and increase the value of this parameter. Suppose that the value of the configuration item is A, total HDFS storage space is B, the threshold is C, and HDFS space used by Hive is D. Adjust the value of the configuration item according to  $A \times B \times C > D$ . The total HDFS storage space can be viewed on the HDFS monitoring page, and HDFS space used by Hive can be viewed on the Hive monitoring page.

2. Check whether the alarm is cleared.
  - If yes, no further action is required.

- If no, go to [Step 2.1](#).

**Step 2** Expand the system.

1. Add nodes.
2. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3.1](#).

**Step 3** Check whether the data node is normal.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether ALM-12006 Node Fault, ALM-12007 Process Fault, or ALM-14002 DataNode Disk Usage Exceeds the Threshold exists.
  - If yes, go to [Step 3.3](#).
  - If no, go to [Step 4](#).
3. Clear the alarm by following the steps provided in ALM-12006 Node Fault, ALM-12007 Process Fault, or ALM-14002 DataNode Disk Usage Exceeds the Threshold.
4. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.468 ALM-16002 Hive SQL Execution Success Rate Is Lower Than the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the percentage of the HiveQL statements that are executed successfully every 30 seconds. Percentage of HiveQL statements that are executed successfully = Number of HiveQL statements that are executed successfully by Hive in a specified period/Total number of HiveQL statements that are executed by Hive. This indicator can be viewed on the Hive service monitoring page. This alarm is generated when the percentage of the HiveQL statements that are executed successfully exceeds the specified threshold (90% by default). The name of the host for which the alarm is generated can be obtained from the location information of the alarm. The host IP address is the IP address of the HiveServer node.

This alarm is cleared when the percentage of the HiveQL statements that are executed successfully in a test period is less than or equal to the threshold.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16002    | Major          | Yes        |

## Parameters

| Parameter         | Description                                             |
|-------------------|---------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated. |
| RoleName          | Specifies the role for which the alarm is generated.    |
| HostName          | Specifies the host for which the alarm is generated.    |
| Trigger condition | Specifies the threshold for triggering the alarm.       |

## Impact on the System

The system configuration and performance cannot meet service processing requirements.

## Possible Causes

- A syntax error occurs in HiveQL commands.
- The HBase service is abnormal when a Hive on HBase task is being performed.
- Basic services that are depended on are abnormal, such as HDFS, Yarn, and ZooKeeper.

## Procedure

**Step 1** Check whether the HiveQL commands comply with syntax.

1. Use the Hive client to log in to the HiveServer node for which the alarm is generated. Query the HiveQL syntax standard provided by Apache, and check whether the HiveQL commands are correct. For details, see <https://cwiki.apache.org/confluence/display/hive/languagemanual>.
  - If yes, go to **Step 2.1**.
  - If no, go to **Step 1.2**.

 **NOTE**

To view the user who runs an incorrect statement, download HiveServerAudit logs of the HiveServer node for which this alarm is generated. Set **Start time** and **End time** to 10 minutes before and after the alarm generation time respectively. Open the log file and search for the **Result=FAIL** keyword to filter the log information about the incorrect statement, and then view the user who runs the incorrect statement according to **UserName** in the log information.

2. Enter correct HiveQL statements, and check whether the command can be properly executed.
  - If yes, go to [Step 4.5](#).
  - If no, go to [Step 2.1](#).

**Step 2** Check whether the HBase service is abnormal.

1. Check whether a Hive on HBase task is performed.
  - If yes, go to [Step 2.2](#).
  - If no, go to [Step 3.1](#).
2. Check whether the HBase service is normal in the service list.
  - If yes, go to [Step 3.1](#).
  - If no, go to [Step 2.3](#).
3. Check the alarms displayed on the alarm page and clear them according to **Alarm Help**.
4. Enter correct HiveQL statements, and check whether the command can be properly executed.
  - If yes, go to [Step 4.5](#).
  - If no, go to [Step 3.1](#).

**Step 3** Check whether the Spark service is abnormal.

1. Check whether the Spark service is normal in the service list.
  - If yes, go to [Step 4.1](#).
  - If no, go to [Step 3.2](#).
2. Check the alarms displayed on the alarm page and clear them according to **Alarm Help**.
3. Enter correct HiveQL statements, and check whether the command can be properly executed.
  - If yes, go to [Step 4.5](#).
  - If no, go to [Step 4.1](#).

**Step 4** Check whether HDFS, Yarn, and ZooKeeper are normal.

1. Go to the MRS cluster details page and click **Components**.
2. In the service list, check whether the services, such as HDFS, Yarn, and ZooKeeper are normal.
  - If yes, go to [Step 4.5](#).
  - If no, go to [Step 4.3](#).
3. Check the alarms displayed on the alarm page and clear them according to **Alarm Help**.

4. Enter correct HiveQL statements, and check whether the command can be properly executed.
  - If yes, go to [Step 4.5](#).
  - If no, go to [Step 5](#).
5. Wait one minute and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 5](#).

**Step 5** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.469 ALM-16004 Hive Service Unavailable (For MRS 2.x or Earlier)

### Description

The system checks the Hive service status every 30 seconds. This alarm is generated when the Hive service is unavailable.

This alarm is cleared when the Hive service recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16004    | Critical       | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

The system cannot provide data loading, query, and extraction services.

## Possible Causes

- Basic services, such as ZooKeeper, HDFS, Yarn, and DBService work incorrectly, or the Hive process is faulty.
  - ZooKeeper is abnormal.
  - HDFS is abnormal.
  - Yarn is abnormal.
  - DBService is abnormal.
  - The Hive service process is faulty. If the alarm is caused by a Hive process fault, the alarm report has a delay of about 5 minutes.
- The network communication between the Hive service and basic services is interrupted.

## Procedure

**Step 1** Check the HiveServer/MetaStore process status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **Hive > Instances**. In the Hive instance list, check whether the status of all HiveServer/MetaStore instances is **Unknown**.
  - If yes, go to [Step 1.3](#).
  - If no, go to [Step 2](#).
3. Above the Hive instance list, choose **More > Restart Instance** to restart the HiveServer/MetaStore process.
4. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Check the ZooKeeper status.

1. Go to the cluster details page and choose **Alarms**.
2. On MRS Manager, check whether the ALM-12007 Process Fault alarm is reported.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 3](#).
3. In the **Alarm Details** area of ALM-12007 Process Fault, check whether **ServiceName** is **ZooKeeper**.
  - If yes, go to [Step 2.4](#).
  - If no, go to [Step 3](#).
4. Rectify the fault by following steps provided in ALM-12007 Process Fault.
5. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.



- If yes, no further action is required.
- If no, go to [Step 3](#).

**Step 3** Check the HDFS status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, check whether the alarm ALM-14000 HDFS Service Unavailable exists.
  - If yes, go to [Step 3.3](#).
  - If no, go to [Step 4](#).
3. Rectify the fault by following the steps provided in ALM-14000 HDFS Service Unavailable.
4. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 4** Check the Yarn status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list on MRS Manager, check whether the alarm ALM-18000 Yarn Service Unavailable is generated.
  - If yes, go to [Step 4.3](#).
  - If no, go to [Step 4](#).
3. Rectify the fault by following the steps provided in ALM-18000 Yarn Service Unavailable.
4. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 5** Check the DBService status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list on MRS Manager, check whether ALM-27001 DBService Unavailable is generated.
  - If yes, go to [Step 5.3](#).
  - If no, go to [Step 6](#).
3. Rectify the fault by following the handling procedure in [ALM-27001 DBService Unavailable \(For MRS 2.x or Earlier\)](#).
4. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 6](#).

**Step 6** Check the network connection between Hive and ZooKeeper, HDFS, Yarn, and DBService.

1. Go to the MRS cluster details page and click **Components**.

2. Click **Hive**.
3. Click **Instances**.  
The HiveServer instance list is displayed.
4. Click **Host Name** in the row of **HiveServer**.  
The HiveServer host status page is displayed.
5. Record the IP address under **Summary**.
6. Use the IP address obtained in [Step 6.5](#) to log in to the host where HiveServer is located.
7. Run the **ping** command to check whether the network connection between the host that runs HiveServer and the hosts that run the ZooKeeper, HDFS, Yarn, and DBService services is normal. Methods of obtaining IP addresses of the hosts that run ZooKeeper, HDFS, Yarn, and DBService services as well as the HiveServer IP address are the same.
  - If yes, go to [Step 7](#).
  - If no, go to [Step 6.8](#).
8. Contact the O&M personnel to restore the network.
9. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 7](#).

**Step 7** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.470 ALM-16005 Number of Failed Hive SQL Executions in the Last Period Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks whether the number of Hive SQL statements that fail to be executed has exceeded the threshold in the last 10-minute period. This alarm is generated when the number of failed Hive SQL statement executions in the last 10 minutes is greater than the threshold. In the next 10 minutes, if the number of failed Hive SQL statement executions is less than the threshold, the alarm is automatically cleared.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 16005    | Major          | Yes        |

## Parameter

| Parameter   | Description                               |
|-------------|-------------------------------------------|
| ServiceName | Service for which the alarm is generated. |
| RoleName    | Role for which the alarm is generated.    |
| HostName    | Host for which the alarm is generated.    |

## Impact on the System

None

## Possible Causes

The Hive SQL syntax is incorrect. As a result, the Hive SQL statements fail to be executed.

## Procedure

Check the Hive SQL statements that fail to be executed, correct the syntax, and execute the SQL statements again.

## Reference

None

## 7.12.471 ALM-18000 Yarn Service Unavailable (For MRS 2.x or Earlier)

### Description

The alarm module checks the Yarn service status every 30 seconds. This alarm is generated when the Yarn service is unavailable.

This alarm is cleared when the Yarn service recovers.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18000    | Critical       | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

The cluster cannot provide the Yarn service. Users cannot run new applications. Submitted applications cannot be run.

## Possible Causes

- ZooKeeper is abnormal.
- HDFS is abnormal.
- There is no active ResourceManager node in the Yarn cluster.
- All NodeManager nodes in the Yarn cluster are abnormal.

## Procedure

**Step 1** Check the ZooKeeper status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, check whether the alarm ALM-13000 ZooKeeper Service Unavailable exists.
  - If yes, go to **Step 1.3**.
  - If no, go to **Step 2.2**.
3. Rectify the fault by following the handling procedure in **ALM-13000 ZooKeeper Service Unavailable (For MRS 2.x or Earlier)**. Then, check whether this alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 2.2**.

**Step 2** Check the HDFS status.

1. Go to the cluster details page and choose **Alarms**.

2. In the alarm list, check whether an HDFS alarm is generated.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 3.2](#).
3. Click **Alarms**, and handle HDFS alarms according to **Alarm Help**. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3.2](#).

**Step 3** Check the ResourceManager status in the Yarn cluster.

1. Go to the MRS cluster details page and click **Components**.
2. Click **Yarn**.
3. In **Yarn Summary**, check whether there is an active ResourceManager node in the Yarn cluster.
  - If yes, go to [Step 4.2](#).
  - If no, go to [Step 5](#).

**Step 4** Check the NodeManager node status in the Yarn cluster.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **Yarn > Instances**.
3. Check **Health Status** of NodeManager, and check whether there are unhealthy nodes.
  - If yes, go to [Step 4.4](#).
  - If no, go to [Step 5](#).
4. Rectify the fault by following the procedure provided in [ALM-18002 NodeManager Heartbeat Lost \(For MRS 2.x or Earlier\)](#) or [ALM-18003 NodeManager Unhealthy \(For MRS 2.x or Earlier\)](#). Then, check whether this alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 5](#).

**Step 5** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.472 ALM-18002 NodeManager Heartbeat Lost (For MRS 2.x or Earlier)

### Description

The system checks the number of lost NodeManager nodes every 30 seconds, and compares the number of lost nodes with the threshold. The **Lost Nodes** indicator

has a default threshold. This alarm is generated when the value of the **Lost Nodes** indicator exceeds the threshold.

This alarm is cleared when the value of **Lost Nodes** is less than or equal to the threshold.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18002    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

- The lost NodeManager node cannot provide the Yarn service.
- The number of containers decreases, so the cluster performance deteriorates.

## Possible Causes

- NodeManager is forcibly deleted without decommission.
- All NodeManager instances are stopped or the NodeManager process is faulty.
- The host where the NodeManager node resides is faulty.
- The network between the NodeManager and ResourceManager is disconnected or busy.

## Procedure

**Step 1** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.473 ALM-18003 NodeManager Unhealthy (For MRS 2.x or Earlier)

### Description

The system checks the number of abnormal NodeManager nodes every 30 seconds, and compares the number of abnormal nodes with the threshold. The **Unhealthy Nodes** indicator has a default threshold. This alarm is generated when the value of the **Unhealthy Nodes** indicator exceeds the threshold.

This alarm is cleared when the value of **Unhealthy Nodes** is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18003    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

### Impact on the System

- The faulty NodeManager node cannot provide the Yarn service.
- The number of containers decreases, so the cluster performance deteriorates.

### Possible Causes

- The disk space of the host where the NodeManager node resides is insufficient.

- User **omm** does not have the permission to access a local directory on the NodeManager node.

## Procedure

**Step 1** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

# 7.12.474 ALM-18004 NodeManager Disk Usability Ratio Is Lower Than the Threshold (For MRS 2.x or Earlier)

## Description

The system checks the available disk space of each NodeManager node every 30 seconds and compares the disk availability rate with the threshold. A default threshold range is provided for the **NodeManager Disk Usability Ratio**. This alarm is generated when the system detects that the actual **NodeManager Disk Usability Ratio** is lower than the threshold.

This alarm is automatically cleared when the value of **NodeManager Disk Usability Ratio** is greater than the threshold.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18004    | Major          | Yes        |

## Parameters

| Parameter         | Description                                             |
|-------------------|---------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated. |
| RoleName          | Specifies the role for which the alarm is generated.    |
| HostName          | Specifies the host for which the alarm is generated.    |
| Trigger condition | Specifies the threshold for triggering the alarm.       |



## Impact on the System

- The NodeManager node whose disk availability rate is lower than the threshold may fail to provide the Yarn service.
- The number of containers decreases, so the cluster performance may deteriorate.

## Possible Causes

- The disk space of the host where the NodeManager node resides is insufficient.
- User **omm** does not have the permission to access a local directory on the NodeManager node.

## Procedure

**Step 1** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.475 ALM-18006 MapReduce Job Execution Timeout (For MRS 2.x or Earlier)

### Description

The alarm module checks the MapReduce job execution every 30 seconds. This alarm is generated when the execution of a submitted MapReduce job times out.

This alarm must be manually cleared.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18006    | Major          | No         |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

Execution of the submitted MapReduce job times out, so no execution result can be obtained. Execute the job again after rectifying the fault.

## Possible Causes

It takes a long time to execute a MapReduce job. However, the specified time is less than the required execution time.


## Procedure

**Step 1** Check whether time is improperly set.

Set **-Dapplication.timeout.interval** to a larger value, or do not set the parameter. Check whether the MapReduce job can be executed.

- If yes, go to [Step 2.5](#).
- If no, go to [Step 2.2](#).

**Step 2** Check the Yarn status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list on MRS Manager, check whether the alarm ALM-18000 Yarn Service Unavailable is generated.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 3](#).
3. Rectify the fault by following the handling procedure in [ALM-18000 Yarn Service Unavailable \(For MRS 2.x or Earlier\)](#).
4. Run the MapReduce job command again to check whether the MapReduce job can be executed.
  - If yes, go to [Step 2.5](#).
  - If no, go to [Step 4](#).
5. In the alarm list, click  in the **Operation** column of the alarm to manually clear the alarm. No further action is required.

**Step 3** Adjust the timeout threshold.

On MRS Manager, choose **System > Threshold Configuration > Services > Yarn > Timed out Applications**, and increase the maximum number of timeout tasks allowed by the current threshold rule. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.476 ALM-18008 Heap Memory Usage of Yarn ResourceManager Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the heap memory usage of Yarn ResourceManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Yarn ResourceManager exceeds the threshold (80% of the maximum memory by default).

To change the threshold, choose **System > Threshold Configuration > Service > Yarn**. The alarm is cleared when the heap memory usage is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18008    | Major          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

| Parameter         | Description                                       |
|-------------------|---------------------------------------------------|
| Trigger Condition | Specifies the threshold for triggering the alarm. |

## Impact on the System

When the heap memory usage of Yarn ResourceManager is overhigh, the performance of Yarn task submission and operation is affected. What is more, a memory overflow occurs so that the Yarn service is unavailable.

## Possible Causes

The heap memory of the Yarn ResourceManager instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

## Procedure

**Step 1** Check the heap memory usage.

1. Go to the MRS cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **18008** and view the IP address and role name of the instance in **Location**.
3. Choose **Components > Yarn > Instances > ResourceManager** (IP address of the instance for which the alarm is generated) > **Customize > Percentage of Used Heap Memory of the ResourceManager**. Check the heap memory usage.
4. Check whether the heap memory usage of ResourceManager has reached the threshold (80% of the maximum memory).
  - If yes, go to **Step 1.5**.
  - If no, go to **Step 2**.
5. Choose **Components > Yarn > Service Configuration**. Set **Type** to **All** and choose **ResourceManager > System**. Change the values of **-Xmx** and **-Xms** in the **GC\_OPTS** parameter based on the site requirements to ensure that the value of **-Xms** is less than that of **-Xmx**. Click **Save Configuration** and select **Restart Role Instance**. Click **OK**.
6. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 2**.

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

### 7.12.477 ALM-18009 Heap Memory Usage of MapReduce JobHistoryServer Exceeds the Threshold (For MRS 2.x or Earlier)

#### Description

The system checks the heap memory usage of MapReduce JobHistoryServer every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of MapReduce JobHistoryServer exceeds the threshold (80% of the maximum memory by default).

To change the threshold, choose **System > Threshold Configuration > Service > MapReduce**. The alarm is cleared when the heap memory usage is less than or equal to the threshold.

#### Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18009    | Major          | Yes                   |

#### Parameters

| Parameter         | Description                                             |
|-------------------|---------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated. |
| RoleName          | Specifies the role for which the alarm is generated.    |
| HostName          | Specifies the host for which the alarm is generated.    |
| Trigger Condition | Specifies the threshold for triggering the alarm.       |

#### Impact on the System

When the heap memory usage of MapReduce JobHistoryServer is overhigh, the performance of MapReduce log archiving is affected. What is more, a memory overflow occurs so that the Yarn service is unavailable.

## Possible Causes

The heap memory of the MapReduce JobHistoryServer instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

## Procedure

### Step 1 Check the heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **18009** and view the IP address and role name of the instance in **Location**.
3. Choose **Components > MapReduce > Instance > JobHistoryServer** (IP address of the instance for which the alarm is generated) > **Customize > JobHistoryServer Heap Memory Usage Statistics**. Check the heap memory usage.
4. Check whether the heap memory usage of JobHistoryServer has reached the threshold (80% of the maximum heap memory).
  - If yes, go to **Step 1.5**.
  - If no, go to **Step 2**.
5. Choose **Components > MapReduce > Service Configuration**. Set **Type** to **All** and choose **JobHistoryServer > System**. Increase the value of **-Xmx** in the **GC\_OPTS** parameter as required, click **Save Configuration**, and select **Restart the affected services or instances**. Click **OK**.
6. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 2**.

### Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.478 ALM-18010 Number of Pending Yarn Tasks Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the number of pending Yarn tasks every 30 seconds and compares the number of tasks with the threshold. This alarm is generated when the number of pending tasks exceeds the threshold.

You can change the threshold by choosing **System > Configure Alarm Threshold > Service > Yarn > Queue Root Pending Applications > Queue Root Pending Applications** on MRS Manager.

This alarm is cleared when the number of pending tasks is less than or equal to the threshold.

## Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18010    | Major          | Yes                   |

## Parameters

| Parameter         | Description                                             |
|-------------------|---------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated. |
| RoleName          | Specifies the role for which the alarm is generated.    |
| HostName          | Specifies the host for which the alarm is generated.    |
| Trigger Condition | Specifies the threshold for triggering the alarm.       |

## Impact on the System

Tasks may be stacked and cannot be processed in a timely manner.

## Possible Causes

The computing capability of the cluster is lower than the task submission rate. As a result, the task cannot be processed in a timely manner after being submitted.

## Procedure

**Step 1** Check the usage of memory and vCores on the Yarn page.

Check whether the values of **Memory Used|Memory Total** and **VCores Used|VCores Total** on the native Yarn page reach or approach the maximum values.

- If yes, go to [Step 2](#).
- If no, go to [Step 5](#).

**Step 2** Check the number of submitted tasks.

Check whether the running tasks are submitted at a normal frequency.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

**Step 3** Scale out the cluster.

The scale-out is based on the site requirements. For details, see [Manually Scaling Out a Cluster](#).

**Step 4** After the scale-out is completed, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Step 5** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.479 ALM-18011 Memory of Pending Yarn Tasks Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the memory of pending Yarn tasks every 30 seconds and compares the memory with the threshold. This alarm is generated when the memory of pending tasks exceeds the threshold.

You can change the threshold by choosing **System > Configure Alarm Threshold > Service > Yarn > Queue Root Pending Memory > Queue Root Pending Memory** on MRS Manager.

This alarm is cleared when the memory of pending tasks is less than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 18011    | Major          | Yes                   |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |



| Parameter         | Description                                          |
|-------------------|------------------------------------------------------|
| HostName          | Specifies the host for which the alarm is generated. |
| Trigger Condition | Specifies the threshold for triggering the alarm.    |

## Impact on the System

Tasks may be stacked and cannot be processed in a timely manner.

## Possible Causes

The computing capability of the cluster is lower than the task submission rate. As a result, the task cannot be processed in a timely manner after being submitted.

## Procedure

**Step 1** Check the usage of memory and vCores on the Yarn page.

Check whether the values of **Memory Used|Memory Total** and **VCores Used|VCores Total** on the native Yarn page reach or approach the maximum values.

- If yes, go to [Step 2](#).
- If no, go to [Step 5](#).

**Step 2** Check the number of submitted tasks.

Check whether the running tasks are submitted at a normal frequency.

- If yes, go to [Step 3](#).
- If no, go to [Step 5](#).

**Step 3** Scale out the cluster.

The scale-out is based on the site requirements. For details, see [Manually Scaling Out a Cluster](#).

**Step 4** After the scale-out is completed, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Step 5** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.480 ALM-18012 Number of Terminated Yarn Tasks in the Last Period Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the number of terminated Yarn tasks every 10 minutes. This alarm is generated when the number of terminated Yarn tasks in the last 10 minutes is greater than the threshold. This alarm is automatically cleared when the number of terminated Yarn tasks is less than the threshold in the next 10 minutes.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18012    | Major          | Yes        |

### Parameter

| Parameter   | Description                               |
|-------------|-------------------------------------------|
| ServiceName | Service for which the alarm is generated. |
| RoleName    | Role for which the alarm is generated.    |
| HostName    | Host for which the alarm is generated.    |

### Impact on the System

None

### Possible Causes

A user manually stops a running Yarn task.

### Procedure

Check the task termination operator in the Yarn logs and audit logs, and determine the cause of the task termination.

### Reference

None

## 7.12.481 ALM-18013 Number of Failed Yarn Tasks in the Last Period Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the number of failed Yarn tasks every 10 minutes. This alarm is generated when the number of failed Yarn tasks in the last 10 minutes is greater than the threshold. This alarm is automatically cleared when the number of failed Yarn tasks is less than the threshold in the next 10 minutes.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 18013    | Major          | Yes        |

### Parameter

| Parameter   | Description                               |
|-------------|-------------------------------------------|
| ServiceName | Service for which the alarm is generated. |
| RoleName    | Role for which the alarm is generated.    |
| HostName    | Host for which the alarm is generated.    |

### Impact on the System

None

### Possible Causes

The submitted Yarn job program is incorrect. For example, the parameter for Spark to submit a job is incorrect.

### Procedure

Check the log of the failed job, locate the failure cause, modify the job, and submit the job again.

### Reference

None

## 7.12.482 ALM-19000 HBase Service Unavailable (For MRS 2.x or Earlier)

### Description

The alarm module checks the HBase service status every 30 seconds. This alarm is generated when the HBase service is unavailable.

This alarm is cleared when the HBase service recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 19000    | Critical       | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

### Impact on the System

Operations cannot be performed, such as reading or writing data and creating tables.

### Possible Causes

- ZooKeeper is abnormal.
- HDFS is abnormal.
- HBase is abnormal.
- The network is abnormal.

### Procedure

**Step 1** Check the ZooKeeper status.

1. Go to the MRS cluster details page and click **Components**.
2. In the service list, check whether the health status of ZooKeeper is **Good**.
  - If yes, go to [Step 2.1](#).

- If no, go to [Step 1.3](#).
- 3. In the alarm list, check whether the alarm ALM-13000 ZooKeeper Service Unavailable exists.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 2.1](#).
- 4. Rectify the fault by following the steps provided in ALM-13000 ZooKeeper Service Unavailable.
- 5. Wait several minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

**Step 2** Check the HDFS status.

1. On MRS Manager, check whether the ALM-14000 HDFS Service Unavailable alarm is reported.
  - If yes, go to [Step 2.2](#).
  - If no, go to [Step 3](#).
2. Rectify the fault by following the steps provided in ALM-14000 HDFS Service Unavailable.
3. Wait several minutes and check whether the alarm is cleared.

**Step 3** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.483 ALM-19006 HBase Replication Sync Failed (For MRS 2.x or Earlier)

### Description

This alarm is generated when disaster recovery (DR) data fails to be synchronized to a standby cluster.

This alarm is cleared when DR data synchronization succeeds.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 19006    | Major          | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

HBase data in a cluster fails to be synchronized to the standby cluster, causing data inconsistency between active and standby clusters.

## Possible Causes

- The HBase service on the standby cluster is abnormal.
- The network is abnormal.

## Procedure

**Step 1** Observe whether the system automatically clears the alarm.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, click the alarm to obtain alarm generation time from **Generated Time** in **Alarm Details**. Check whether the alarm has existed for over 5 minutes.
  - If yes, go to [Step 2.1](#).
  - If no, go to [Step 1.3](#).
3. Wait 5 minutes and check whether the alarm is automatically cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

**Step 2** Check the HBase service status of the standby cluster.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, click the alarm and obtain **HostName** from **Location** in **Alarm Details**.
3. Log in to the node where the HBase client of the active cluster is located. Run the following commands to switch the user:  
**sudo su - root**  
**su - omm**
4. Run the **status 'replication', 'source'** command to check the synchronization status of the faulty node.  
The synchronization status of a node is as follows.

```
10-10-10-153:
SOURCE: PeerID=abc, SizeOfLogQueue=0, ShippedBatches=2, ShippedOps=2, ShippedBytes=320,
LogReadInBytes=1636, LogEditsRead=5, LogEditsFiltered=3, SizeOfLogToReplicate=0,
TimeForLogToReplicate=0, ShippedHFiles=0, SizeOfHFileRefsQueue=0, AgeOfLastShippedOp=0,
TimeStampsOfLastShippedOp=Mon Jul 18 09:53:28 CST 2016, Replication Lag=0,
FailedReplicationAttempts=0
SOURCE: PeerID=abc1, SizeOfLogQueue=0, ShippedBatches=1, ShippedOps=1, ShippedBytes=160,
LogReadInBytes=1636, LogEditsRead=5, LogEditsFiltered=3, SizeOfLogToReplicate=0,
TimeForLogToReplicate=0, ShippedHFiles=0, SizeOfHFileRefsQueue=0, AgeOfLastShippedOp=16788,
TimeStampsOfLastShippedOp=Sat Jul 16 13:19:00 CST 2016, Replication Lag=16788,
FailedReplicationAttempts=5
```

5. Obtain **PeerID** corresponding to a record whose **FailedReplicationAttempts** value is greater than 0.

In the preceding step, data on the faulty node **10-10-10-153** fails to be synchronized to a standby cluster whose **PeerID** is **abc1**.

6. Run the **list\_peers** command to find the cluster and the HBase instance corresponding to **PeerID**.

```
PEER_ID CLUSTER_KEY STATE TABLE_CFS
abc1 10.10.10.110,10.10.10.119,10.10.10.133:24002:/hbase2 ENABLED
abc 10.10.10.110,10.10.10.119,10.10.10.133:24002:/hbase ENABLED
```

In the preceding information, **/hbase2** indicates that data is synchronized to the HBase2 instance of the standby cluster.

7. In the service list of the standby cluster, check whether the health status of the HBase instance obtained in [Step 2.6](#) is **Good**.
  - If yes, go to [Step 3.1](#).
  - If no, go to [Step 2.8](#).
8. In the alarm list, check whether the alarm ALM-19000 HBase Service Unavailable exists.
  - If yes, go to [Step 2.9](#).
  - If no, go to [Step 3.1](#).
9. Rectify the fault by following the steps provided in ALM-19000 HBase Service Unavailable.
10. Wait several minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3.1](#).

### Step 3 Check the network connection between RegionServers on active and standby clusters.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, click the alarm and obtain **HostName** from **Location** in **Alarm Details**.
3. Log in to the faulty RegionServer node.
4. Run the **ping** command to check whether the network connection between the faulty RegionServer node and the host where RegionServer of the standby cluster resides is normal.
  - If yes, go to [Step 4](#).
  - If no, go to [Step 3.5](#).
5. Contact the O&M personnel to restore the network.
6. After the network recovers, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.484 ALM-19007 HBase Merge Queue Exceeds the Threshold (for 2.x and Earlier Versions)

### Description

The system checks the HBase compaction queue size every 30 seconds. This alarm is generated when the compaction queue size exceeds the alarm threshold (**100** by default) for three consecutive times. This alarm is cleared when the compaction queue size is less than the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 19007    | Minor          | Yes        |

### Parameters

| Name        | Meaning                                                 |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| Host Name   | Specifies the host for which the alarm is generated.    |

### Impact on the System

The cluster performance may deteriorate, affecting data read and write.



## Possible Causes

- The number of HBase RegionServers is too small.
- There are too many regions on a RegionServer of HBase.
- The HBase RegionServer heap size is small.
- Resources are insufficient.
- Related parameters are not configured properly.

## Procedure

**Step 1** Check whether related HBase parameters are properly configured.

1. Log in to the MRS cluster details page, choose **Components > HBase > Service Configuration**, switch **Basic Configuration** to **All Configurations**, and search for **hbase.hstore.compaction.min** and **hbase.hstore.compaction.max**, and increase the values of **hbase.regionserver.thread.compaction.small** and **hbase.regionserver.thread.compaction.throttle**.

### NOTE

- If you did not synchronize IAM users, perform synchronization first. (In the **Dashboard** tab, click **Synchronize** next to **IAM User Sync**.)
2. Save the configuration, and restart the HBase service during off-peak hours or perform a rolling restart to make the configuration take effect.
3. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

## Related Information

None

## 7.12.485 ALM-20002 Hue Service Unavailable (For MRS 2.x or Earlier)

### Description

The system checks the Hue service status every 60 seconds. This alarm is generated if the Hue service is unavailable.

This alarm is cleared when the Hue service is normal.

## Attribute

| Alarm ID | Alarm Severity | Automatically Cleared |
|----------|----------------|-----------------------|
| 20002    | Critical       | Yes                   |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

The system cannot provide data loading, query, and extraction services.

## Possible Causes

- The KrbServer service on which Hue depends is abnormal.
- The DBService service on which Hue depends is abnormal.
- The network connection to DBService is abnormal.

## Procedure

**Check whether the KrbServer service is normal.**

**Step 1** Go to the MRS cluster details page and click **Components**.

**Step 2** In the service list, check whether **Health Status** of **KrbServer** is **Good**.

- If yes, go to [Step 5](#).
- If no, go to [Step 3](#).

**Step 3** Click **Restart** in the **Operation** column of the KrbServer service to restart the service.

**Step 4** Wait for several minutes. Check whether ALM-20002 Hue Service Unavailable is cleared.

- If yes, no further action is required.
- If no, go to [Step 5](#).

**Check whether DBService is normal.**

**Step 5** Go to the MRS cluster details page and click **Components**.

**Step 6** In the service list, check whether **Health Status** of **DBService** is **Good**.

- If yes, go to [Step 9](#).
- If no, go to [Step 7](#).

**Step 7** Click **Restart** in the **Operation** column of the **DBService** service to restart the service.

 **NOTE**

To restart the service, you need to enter the password of the MRS Manager administrator and select **Start or restart related services**.

**Step 8** Wait for several minutes. Check whether **ALM-20002 Hue Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 9](#).

**Check whether the network connected to DBService is normal.**

**Step 9** Choose **Components > Hue > Instance** and record the IP address of the active Hue node.

**Step 10** Use PuTTY to log in to the active Hue.

**Step 11** Run the **ping** command to check whether the network connection between the host where the active Hue is located and the host where **DBService** is located is normal. (The method of obtaining the **DBService** service IP address is the same as that of obtaining the active Hue IP address.)

- If yes, go to [Step 17](#).
- If no, go to [Step 12](#).

**Step 12** Contact the network administrator to repair the network.

**Step 13** Wait for several minutes. Check whether **ALM-20002 Hue Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to [Step 17](#).

**Collect fault information.**

**Step 14** On MRS Manager, choose **System > Export Log**.

**Step 15** Select the following nodes from the **Services** drop-down list and click **OK**.

- Hue
- Controller

**Step 16** Set **Start Time** and **End Time** for log collection to 10 minutes before and after the alarm is generated, select an export type, and click **OK** to collect the corresponding fault log information.

**Restart Hue.**

**Step 17** Choose **Components > Hue**.

**Step 18** Choose **More > Restart Service** and click **OK**.

**Step 19** Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 20](#).

**Step 20** Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.486 ALM-23001 Loader Service Unavailable (For MRS 2.x or Earlier)

### Description

The system checks the Loader service availability every 60 seconds. This alarm is generated if the Loader service is unavailable and is cleared after the Loader service recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 23001    | Critical       | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

### Impact on the System

Data loading, import, and conversion are unavailable.

## Possible Causes

- The services that Loader depends on are abnormal.
  - ZooKeeper is abnormal.
  - HDFS is abnormal.
  - DBService is abnormal.
  - Yarn is abnormal.
  - MapReduce is abnormal.
- The network is faulty. Loader cannot communicate with its dependent services.
- Loader is running improperly.

## Procedure

### Step 1 Check the ZooKeeper status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **ZooKeeper** and check whether the health status of ZooKeeper is normal.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 1.3](#).
3. Choose **More > Restart Service** to restart ZooKeeper. After ZooKeeper starts, check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 1.4](#).
4. On MRS Manager, check whether the ALM-12007 Process Fault alarm is reported.
  - If yes, go to [Step 1.5](#).
  - If no, go to [Step 2.1](#).
5. In **Alarm Details** of the "ALM-12007 Process Fault" alarm, check whether **ServiceName** is **ZooKeeper**.
  - If yes, go to [Step 1.6](#).
  - If no, go to [Step 2.1](#).
6. Clear the alarm according to the handling suggestions of "ALM-12007 Process Fault".
7. Check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

### Step 2 Check the HDFS status.

1. Go to the MRS cluster details page and choose **Alarms**.
2. On MRS Manager, check whether the "ALM-14000 HDFS Service Unavailable alarm" is reported.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 3.1](#).

3. Clear the alarm according to the handling suggestions of "ALM-14000 HDFS Service Unavailable".
4. Check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3.1](#).

**Step 3** Check the DBService status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **DBService** to check whether the health status of DBService is normal.
  - If yes, go to [Step 4.1](#).
  - If no, go to [Step 3.3](#).
3. Choose **More > Restart Service** to restart DBService. After DBService starts, check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4.1](#).

**Step 4** Check the MapReduce status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **MapReduce** and check whether the health status of MapReduce is normal.
  - If yes, go to [Step 5.1](#).
  - If no, go to [Step 4.3](#).
3. Choose **More > Restart Service** to restart MapReduce. After MapReduce starts, check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 5.1](#).

**Step 5** Check the Yarn status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **Yarn** and check whether the health status of Yarn is normal.
  - If yes, go to [Step 5.4](#).
  - If no, go to [Step 5.3](#).
3. Choose **More > Restart Service** to restart Yarn. After Yarn starts, check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 5.4](#).
4. On MRS Manager, check whether the "ALM-18000 Yarn Service Unavailable" alarm is reported.
  - If yes, go to [Step 5.5](#).
  - If no, go to [Step 6.1](#).
5. Clear the alarm according to the handling suggestions of "ALM-18000 Yarn Service Unavailable".
6. Check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 6.1](#).

**Step 6** Check the network connections between Loader and its dependent components.

1. Go to the MRS cluster details page and click **Components**.
2. Click **Loader**.
3. Click **Instance**. The Sqoop instance list is displayed.
4. Record the management IP addresses of all Sqoop instances.
5. Log in to the hosts using the IP addresses obtained in [Step 6.4](#). Run the following commands to switch the user:  
**sudo su - root**  
**su - omm**
6. Run the **ping** command to check whether the network connection between the hosts where the Sqoop instances reside and the dependent components is normal. (The dependent components include ZooKeeper, DBService, HDFS, MapReduce, and Yarn. The method to obtain the IP addresses of the dependent components is the same as that used to obtain the IP addresses of the Sqoop instances.)
  - If yes, go to [Step 7](#).
  - If no, go to [Step 6.7](#).
7. Contact the network administrator to repair the network.
8. Check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 7](#).

**Step 7** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## 7.12.487 ALM-24000 Flume Service Unavailable (For MRS 2.x or Earlier)

### Description

The alarm module checks the Flume service status every 180 seconds. This alarm is generated if the Flume service is abnormal.

This alarm is cleared after the Flume service recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 24000    | Critical       | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

Flume cannot work and data transmission is interrupted.

## Possible Causes

- HDFS is unavailable.
- LdapServer is unavailable.

## Procedure

**Step 1** Check the HDFS status.

1. Go to the MRS cluster details page and choose **Alarms**.
2. Check whether the ALM-14000 HDFS Service Unavailable alarm is generated.
  - If yes, clear the alarm according to the handling suggestions of "ALM-14000 HDFS Service Unavailable".
  - If no, go to [Step 2](#).

**Step 2** Check the LdapServer status.

Check whether the ALM-25000 LdapServer Service Unavailable alarm is generated.

- If yes, clear the alarm according to the handling suggestions of "ALM-25000 LdapServer Service Unavailable".
- If no, go to [Step 3.2](#).

**Step 3** Check whether the HDFS and LdapServer services are stopped.

1. Go to the MRS cluster details page and click **Components**.
2. In the service list on MRS Manager, check whether the HDFS and LdapServer services are stopped.
  - If yes, start the HDFS and LdapServer services and go to [Step 3.3](#).
  - If no, go to [Step 4](#).
3. Check whether the "ALM-24000 Flume Service Unavailable" alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 4** Collect fault information.



1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.488 ALM-24001 Flume Agent Is Abnormal (For MRS 2.x or Earlier)

### Description

This alarm is generated if the Flume agent monitoring module detects that the Flume agent process is abnormal.

This alarm is cleared after the Flume agent process recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 24001    | Minor          | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

### Impact on the System

Functions of the alarmed Flume agent instance are abnormal. Data transmission tasks of the instance are suspended. In real-time data transmission, data will be lost.

### Possible Causes

- The **JAVA\_HOME** directory does not exist or the Java permission is incorrect.
- The permission of the Flume agent directory is incorrect.

## Procedure

### Step 1 Check the Flume agent's configuration file.

1. Log in to the host where the faulty node resides. Run the following command to switch to user **root**:  
**sudo su - root**
2. Run the **cd *Flume installation directory*/fusioninsight-flume-1.6.0/conf/** command to go to Flume's configuration directory.
3. Run the **cat ENV\_VARS** command. Check whether the **JAVA\_HOME** directory exists and whether the Flume agent user has execute permission of Java.
  - If yes, go to [Step 2.1](#).
  - If no, go to [Step 1.4](#).
4. Specify the correct JAVA\_HOME directory and grant the Flume agent user with the execute permission of Java. Then go to [Step 2.4](#).

### Step 2 Check the permission of the Flume agent directory.

1. Log in to the host where the faulty node resides. Run the following command to switch to user **root**:  
**sudo su - root**
2. Run the following command to access the installation directory of the Flume agent:  
**cd *Flume agent installation directory***
3. Run the **ls -al \* -R** command. Check whether the owner of all files is the Flume agent user.
  - If yes, go to [Step 3](#).
  - If no, run the **chown** command and change the owner of the files to the Flume agent user. Then go to [Step 2.4](#).
4. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3](#).

### Step 3 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.489 ALM-24003 Flume Client Connection Interrupted (For MRS 2.x or Earlier)

### Description

The alarm module monitors the port connection status on the Flume server. This alarm is generated if the Flume server fails to receive a connection message from the Flume client in 3 consecutive minutes.

This alarm is cleared after the Flume server receives a connection message from the Flume client.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 24003    | Major          | Yes        |

### Parameters

| Parameter  | Description                                   |
|------------|-----------------------------------------------|
| ClientIP   | Specifies the IP address of the Flume client. |
| ServerIP   | Specifies the IP address of the Flume server. |
| ServerPort | Specifies the port on the Flume server.       |

### Impact on the System

The communication between the Flume client and server fails. The Flume client cannot send data to the Flume server.

### Possible Causes

- The network between the Flume client and server is faulty.
- The Flume client's process is abnormal.
- The Flume client is incorrectly configured.

### Procedure

**Step 1** Check the network between the Flume client and server.

1. Log in to the host where the alarmed Flume client resides. Run the following command to switch to user **root**:

```
sudo su - root
```

2. Run the **ping** *Flume server IP address* command to check whether the network between the Flume client and server is normal.
  - If yes, go to [Step 2.1](#).
  - If no, go to [Step 4](#).

**Step 2** Check whether the Flume client's process is normal.

1. Log in to the host where the alarmed Flume client resides. Run the following command to switch to user **root**:  
**sudo su - root**
2. Run the **ps -ef|grep flume |grep client** command to check whether the Flume client process exists.
  - If yes, go to [Step 3.1](#).
  - If no, go to [Step 4](#).

**Step 3** Check the Flume client configuration.

1. Log in to the host where the alarmed Flume client resides. Run the following command to switch to user **root**:  
**sudo su - root**
2. Run the **cd** *Flume installation directory/fusioninsight-flume-1.6.0/conf/* command to go to Flume's configuration directory.
3. Run the **cat properties.properties** command to query the current configuration file of the Flume client.
4. Check whether the **properties.properties** file is correctly configured according to the configuration description of the Flume agent.
  - If yes, go to [Step 3.5](#).
  - If no, go to [Step 4](#).
5. Modify the **properties.properties** configuration file.
6. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.490 ALM-24004 Flume Fails to Read Data (For MRS 2.x or Earlier)

### Description

The alarm module monitors the Flume source status. This alarm is generated if the duration that Flume source fails to read data exceeds the threshold.

Users can modify the threshold as required.

This alarm is cleared if the source reads data successfully.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 24004    | Major          | Yes        |

### Parameters

| Parameter     | Description                                                    |
|---------------|----------------------------------------------------------------|
| ServiceName   | Specifies the service for which the alarm is generated.        |
| HostName      | Specifies the host for which the alarm is generated.           |
| ComponentType | Specifies the component type for which the alarm is generated. |
| ComponentName | Specifies the component name for which the alarm is generated. |

### Impact on the System

Data collection is stopped.

### Possible Causes

- The Flume source is faulty.
- The network is faulty.

### Procedure

**Step 1** Check whether the Flume source is normal.

1. Check whether the Flume source is the spoolDir type.
  - If yes, go to [Step 1.2](#).

- If no, go to [Step 1.3](#).
- 2. Query the **spoolDir** directory and check whether all files have been sent.
  - If yes, no further action is required.
  - If no, go to [Step 1.5](#).
- 3. Check whether the Flume source is the Kafka type.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 1.5](#).
- 4. Log in to the Kafka client and run the following commands to check whether all topic data configured for the Kafka source has been consumed.  
**cd /opt/client/Kafka/kafka/bin**  
**./kafka-consumer-groups.sh --bootstrap-server *Kafka cluster IP address:21007* --new-consumer --describe --group *example-group1* --command-config**  
**../config/consumer.properties**
  - If yes, no further action is required.
  - If no, go to [Step 1.5](#).
- 5. Go to the cluster details page and click **Components**.
- 6. Choose **Flume > Instances**.
- 7. Click the Flume instance of the faulty node and check whether the value of the **Source Speed Metrics** is 0.
  - If yes, go to [Step 2.1](#).
  - If no, no further action is required.

**Step 2** Check the status of the network between the Flume source and faulty node.

1. Check whether the Flume source is the avro type.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 3](#).
2. Log in to the host where the faulty node resides. Run the following command to switch to user **root**:  
**sudo su - root**
3. Run the **ping *Flume source IP address*** command to check whether the Flume source can be pinged.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 2.4](#).
4. Contact the network administrator to repair the network.
5. Wait for a while and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3](#).

**Step 3** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.491 ALM-24005 Data Transmission by Flume Is Abnormal (For MRS 2.x or Earlier)

### Description

The alarm module monitors the capacity of Flume channels. This alarm is generated if the duration that a channel is full or the number of times that a source fails to send data to the channel exceeds the threshold.

Users can set the threshold as required by modifying the **channelfullcount** parameter.

This alarm is cleared after the Flume channel space is released.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 24005    | Major          | Yes        |

### Parameters

| Parameter     | Description                                                    |
|---------------|----------------------------------------------------------------|
| ServiceName   | Specifies the service for which the alarm is generated.        |
| HostName      | Specifies the host for which the alarm is generated.           |
| ComponentType | Specifies the component type for which the alarm is generated. |
| ComponentName | Specifies the component name for which the alarm is generated. |

### Impact on the System

If the usage of the Flume channel continues to grow, the data transmission time increases. When the usage reaches 100%, the Flume agent process is suspended.

### Possible Causes

- The Flume sink is faulty.
- The network is faulty.

## Procedure

### Step 1 Check whether the Flume sink is normal.

1. Check whether the Flume sink is the HDFS type.
  - If yes, go to [Step 1.2](#).
  - If no, go to [Step 1.3](#).
2. On MRS Manager, check whether the ALM-14000 HDFS Service Unavailable alarm is reported and whether the HDFS service is stopped.
  - If the alarm is reported, clear it according to the handling suggestions of ALM-14000 HDFS Service Unavailable; if the HDFS service is stopped, start it. Then go to [Step 1.7](#).
  - If no, go to [Step 1.7](#).
3. Check whether the Flume sink is the HBase type.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 1.7](#).
4. On MRS Manager, check whether the ALM-19000 HBase Service Unavailable alarm is reported and whether the HBase service is stopped.
  - If the alarm is reported, clear it according to the handling suggestions of "ALM-19000 HBase Service Unavailable"; if the HBase service is stopped, start it. Then go to [Step 1.7](#).
  - If no, go to [Step 1.7](#).
5. Check whether the Flume sink is the Kafka type.
  - If yes, go to [Step 1.6](#).
  - If no, go to [Step 1.7](#).
6. On MRS Manager, check whether the ALM-38000 Kafka Service Unavailable alarm is reported and whether the Kafka service is stopped.
  - If the alarm is reported, clear it according to the handling suggestions of "ALM-38000 Kafka Service Unavailable"; if the Kafka service is stopped, start it. Then go to [Step 1.7](#).
  - If no, go to [Step 1.7](#).
7. Go to the MRS cluster details page and click **Components**.
8. Choose **Flume > Instances**.
9. Click the Flume instance of the faulty node and check whether the value of the **Sink Speed Metrics** is 0.
  - If yes, go to [Step 2.1](#).
  - If no, no further action is required.

### Step 2 Check the status of the network between the Flume sink and faulty node.

1. Check whether the Flume sink is the Avro type.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 3](#).
2. Log in to the host where the faulty node resides. Run the following command to switch to user **root**:  
**sudo su - root**



3. Run the **ping** *Flume sink IP address* command to check whether the Flume sink can be pinged.
  - If yes, go to [Step 3](#).
  - If no, go to [Step 2.4](#).
4. Contact the network administrator to repair the network.
5. Wait for a while and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3](#).

**Step 3** Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.492 ALM-25000 LdapServer Service Unavailable (For MRS 2.x or Earlier)

### Description

The system checks the LdapServer service status every 30 seconds. This alarm is generated when the active and standby LdapServer services are abnormal.

This alarm is cleared when either of the LdapServer services restores.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 25000    | Critical       | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

When this alarm is generated, no operation can be performed for the KrbServer users and LdapServer users in the cluster. For example, users, user groups, or roles cannot be added, deleted, or modified, and user passwords cannot be changed on MRS Manager. The authentication for existing users in the cluster is not affected.

## Possible Causes

- The node where the LdapServer service locates is faulty.
- The LdapServer process is abnormal.

## Procedure

- Step 1** Check whether the nodes where the two SlapdServer instances of the LdapServer service locate are faulty.
1. Go to the MRS cluster details page and click **Components**.
  2. Choose **LdapServer > Instances**. Go to the LdapServer instance page to obtain the host name of the node where the two SlapdServer instances reside.
  3. On the **Alarms** page of MRS Manager, check whether the alarm ALM-12006 Node Fault is generated.
    - If yes, go to [Step 1.4](#).
    - If no, go to [Step 2.1](#).
  4. Check whether the host name in the alarm information is the same as the actual host name in [Step 1.2](#).
    - If yes, go to [Step 1.5](#).
    - If no, go to [Step 2.1](#).
  5. Rectify the fault by following steps provided in ALM-12006 Node Fault.
  6. In the alarm list, check whether the alarm ALM-25000 LdapServer Service Unavailable is cleared.
    - If yes, no further action is required.
    - If no, go to [Step 3](#).
- Step 2** Check whether the LdapServer process is in normal state.
1. Go to the cluster details page and choose **Alarms**.
  2. Check whether ALM-12007 Process Fault is generated.
    - If yes, go to [Step 2.3](#).
    - If no, go to [Step 3](#).
  3. Check whether the service name and host name in the alarm are consistent with the LdapServer service and host names.
    - If yes, go to [Step 2.4](#).
    - If no, go to [Step 3](#).
  4. Rectify the fault by following steps provided in ALM-12007 Process Fault.
  5. In the alarm list, check whether the alarm ALM-25000 LdapServer Service Unavailable is cleared.
    - If yes, no further action is required.

- If no, go to [Step 3](#).

**Step 3** Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

**Reference**

None

## 7.12.493 ALM-25004 Abnormal LdapServer Data Synchronization (For MRS 2.x or Earlier)

**Description**

This alarm is generated when LdapServer data on Manager is inconsistent. This alarm is cleared when the data becomes consistent.

This alarm is generated when LdapServer data in the cluster is inconsistent with LdapServer data on Manager. This alarm is cleared when the data becomes consistent.

**Attribute**

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 25004    | Critical       | Yes        |

**Parameters**

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Host for which the alarm is generated.                  |

**Impact on the System**

LdapServer data inconsistency occurs because LdapServer data on Manager or in the cluster is damaged. The LdapServer process with damaged data cannot provide services externally, and the authentication functions of Manager and the cluster are affected.

## Possible Causes

- The network of the node where the LdapServer process locates is faulty.
- The LdapServer process is abnormal.
- The OS restart damages data on LdapServer.

## Procedure

**Step 1** Check whether the network where the LdapServer nodes reside is faulty.

1. Go to the cluster details page and choose **Alarms**.
2. Record the IP address of **HostName** in **Location** of the alarm as **IP1** (if multiple alarms exist, record the IP addresses as **IP1**, **IP2**, and **IP3** respectively).
3. Contact O&M personnel and use PuTTY to log in to the node corresponding to **IP1**. Run the **ping** command on the node to check whether the IP address of the management plane of the active OMS node can be pinged.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 2.1](#).
4. Contact O&M personnel to recover the network and check whether the alarm **ALM-25004 Abnormal LdapServer Data Synchronization** is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

**Step 2** Check whether the LdapServer process is in normal state.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether ALM-12004 OLdap Resource Is Abnormal is generated for LdapServer.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 2.5](#).
3. Rectify the fault by following steps provided in **ALM-12004 OLdap Resource Is Abnormal**.
4. Check whether the alarm ALM-25004 Abnormal LdapServer Data Synchronization is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 2.5](#).
5. On the **Alarms** page of MRS Manager, check whether the alarm ALM-12007 Process Fault of LdapServer is generated.
  - If yes, go to [Step 2.6](#).
  - If no, go to [Step 3.1](#).
6. Rectify the fault by following steps provided in ALM-12007 Process Fault.
7. Check whether the alarm ALM-25004 Abnormal LdapServer Data Synchronization is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3.1](#).

**Step 3** Check whether the OS restart damages data on LdapServer.

1. Go to the cluster details page and choose **Alarms**.
2. Record the IP address of **HostName** in **Location** of the alarm as **IP1** (if multiple alarms exist, record the IP addresses as **IP1**, **IP2**, and **IP3** respectively). Choose **Services > LdapServer > Service Configuration** and record the LdapServer port number as **PORT**. (If the IP address in the alarm location information is the IP address of the standby OMS node, the default port number is 21750.)
3. Log in to node **IP1** as user **omm** and run the **ldapsearch -H ldaps://IP1:PORT -x -LLL -b dc=hadoop,dc=com** command (if the IP address is the IP address of the standby OMS node, run the **ldapsearch -H ldaps://IP1:PORT -x -LLL -b dc=hadoop,dc=com** command before running this command). Check whether error information is displayed in the command output.
  - If yes, go to **Step 3.4**.
  - If no, go to **Step 4**.
4. Recover the LdapServer and OMS nodes using backup data before the alarm is generated. For details, see section "Recovering Manager Data" in the *Administrator Guide*.

**NOTE**

Use the OMS data and LdapServer data backed up at the same time to restore data. Otherwise, the service and operation may fail. To recover data when services run properly, you are advised to manually back up the latest management data and then recover the data. Otherwise, Manager data produced between the backup point in time and the recovery point in time will be lost.

5. Check whether the alarm ALM-25004 Abnormal LdapServer Data Synchronization is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 4**.

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

**Reference**

None

**7.12.494 ALM-25500 KrbServer Service Unavailable (For MRS 2.x or Earlier)****Description**

The system checks the KrbServer service status every 30 seconds. This alarm is generated when the KrbServer service is abnormal.

This alarm is cleared when the KrbServer service is in normal state.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 25500    | Critical       | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

When this alarm is generated, no operation can be performed for the KrbServer component in the cluster. The authentication of KrbServer in other components will be affected. The health status of components that depend on KrbServer in the cluster is **Bad**.

## Possible Causes

- The node where the KrbServer service locates is faulty.
- The OLdap service is unavailable.

## Procedure

**Step 1** Check whether the node where the KrbServer service locates is faulty.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **KrbServer > Instances**. Go to the KrbServer instance page and view the host name of the node where the KrbServer service is deployed.
3. On the **Alarms** page of MRS Manager, check whether the alarm ALM-12006 Node Fault is generated.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 2.1](#).
4. Check whether the host name in the alarm information is the same as the actual host name in [Step 1.2](#).
  - If yes, go to [Step 1.5](#).
  - If no, go to [Step 2.1](#).
5. Rectify the fault by following steps provided in ALM-12006 Node Fault.

6. In the alarm list, check whether the alarm ALM-25500 KrbServer Service Unavailable is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3](#).

**Step 2** Check whether the OLdap service is unavailable.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether ALM-12004 OLdap Resource Is Abnormal is generated.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 3](#).
3. Rectify the fault by following steps provided in ALM-12004 OLdap Resource Is Abnormal.
4. In the alarm list, check whether the alarm ALM-25500 KrbServer Service Unavailable is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 3](#).

**Step 3** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.495 ALM-26051 Storm Service Unavailable (For MRS 2.x or Earlier)

### Description

The system checks the Storm service availability every 30 seconds. This alarm is generated if the Storm service becomes unavailable after all Nimbus nodes in a cluster become abnormal.

This alarm is cleared after the Storm service recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 26051    | Critical       | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

- The cluster cannot provide the Storm service.
- Users cannot run new Storm tasks.

## Possible Causes

- The Kerberos component is faulty.
- ZooKeeper is faulty or suspended.
- The active and standby Nimbus nodes in the Storm cluster are abnormal.

## Procedure

**Step 1** Check the Kerberos component status. For clusters without Kerberos authentication, skip this step and go to [Step 2](#).

1. Go to the MRS cluster details page and click **Components**.
2. Check whether the health status of the Kerberos service is **Good**.
  - If yes, go to [Step 2.1](#).
  - If no, go to [Step 1.3](#).
3. Rectify the fault by following instructions in ALM-25500 KrbServer Service Unavailable.
4. Perform [Step 1.2](#) again.

**Step 2** Check the ZooKeeper component status.

1. Check whether the health status of the ZooKeeper service is **Good**.
  - If yes, go to [Step 3.1](#).
  - If no, go to [Step 2.2](#).
2. If the ZooKeeper service is stopped, start it. For other problems, follow the instructions in ALM-13000 ZooKeeper Service Unavailable.
3. Perform [Step 2.1](#) again.

**Step 3** Check the status of the active and standby Nimbus nodes.

1. Choose **Components > Storm > Nimbus**.
2. In **Role**, check whether only one active Nimbus node exists.



- If yes, go to [Step 4](#).
  - If no, go to [Step 3.3](#).
3. Select the two Nimbus instances and choose **More > Restart Instance**. Check whether the restart is successful.
    - If yes, go to [Step 3.4](#).
    - If no, go to [Step 4](#).
  4. Log in to MRS Manager again and choose **Components > Storm > Nimbus**. Check whether the health status of Nimbus is **Good**.
    - If yes, go to [Step 3.5](#).
    - If no, go to [Step 4](#).
  5. Wait 30 seconds and check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.496 ALM-26052 Number of Available Supervisors in Storm Is Lower Than the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the number of supervisors every 60 seconds and compares it with the threshold. This alarm is generated if the number of supervisors is lower than the threshold.

To modify the threshold, users can choose **System > Threshold Configuration** on MRS Manager.

This alarm is cleared if the number of supervisors is greater than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 26052    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger Condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

- Existing tasks in the cluster cannot be executed.
- The cluster can receive new Storm tasks but cannot execute them.

## Possible Causes

Supervisors are abnormal in the cluster.

## Procedure

**Step 1** Check the supervisor status.

1. Go to the cluster details page and click **Components**.
2. Choose **Storm > Supervisor**.
3. In **Role**, check whether the cluster has supervisor instances that are in the **Faulty** or **Recovering** state.
  - If yes, go to **Step 1.4**.
  - If no, go to **Step 2**.
4. Select the supervisor instances that are in the **Faulty** or **Recovering** state and choose **More > Restart Instance**.
  - If yes, go to **Step 1.5**.
  - If the restart fails, go to **Step 2**.
5. Wait 30 seconds and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 2**.

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

# 7.12.497 ALM-26053 Slot Usage of Storm Exceeds the Threshold (For MRS 2.x or Earlier)

## Description

The system checks the slot usage of Storm every 60 seconds and compares it with the threshold. This alarm is generated if the slot usage exceeds the threshold.

To modify the threshold, users can choose **System > Threshold Configuration** on MRS Manager.

This alarm is cleared if the slot usage is lower than or equal to the threshold.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 26053    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System

Users cannot run new Storm tasks.

## Possible Causes

- Supervisors are abnormal in the cluster.
- Supervisors are normal but have poor processing capability.

## Procedure

### Step 1 Check the supervisor status.

1. Go to the cluster details page and click **Components**.
2. Choose **Storm > Supervisor**.
3. In **Role**, check whether the cluster has supervisor instances that are in the **Faulty** or **Recovering** state.
  - If yes, go to **Step 1.4**.
  - If no, go to **Step 2.1** or **Step 3.1**.
4. Select the supervisor instances that are in the **Faulty** or **Recovering** state and choose **More > Restart Instance**.
  - If yes, go to **Step 1.5**.
  - If the restart fails, go to **Step 4**.
5. Wait a moment and then check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 2.1** or **Step 3.1**.

### Step 2 Increase the number of slots for the supervisors.

1. Go to the cluster details page and click **Components**.
2. Choose **Storm > Supervisor > Service Configuration**, and set **Type** to **All**.
3. Increase the value of **supervisor.slots.ports** to increase the number of slots for each supervisor. Then restart the instances.
4. Wait a moment and then check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 4**.

### Step 3 Expand the capacity of the supervisors.

1. Add nodes.
2. Wait a moment and then check whether the alarm is cleared.
  - If yes, no further action is required.
  - If the restart fails, go to **Step 4**.

### Step 4 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.498 ALM-26054 Heap Memory Usage of Storm Nimbus Exceeds the Threshold (For MRS 2.x or Earlier)

### Description

The system checks the heap memory usage of Storm Nimbus every 30 seconds and compares it with the threshold. This alarm is generated if the heap memory usage exceeds the threshold (80% by default).

To modify the threshold, users can choose **System > Threshold Configuration > Service > Storm** on MRS Manager.

This alarm is cleared if the heap memory usage is lower than or equal to the threshold.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 26054    | Major          | Yes        |

### Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| Trigger Condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

### Impact on the System

Frequent memory garbage collection or memory overflow may occur, affecting submission of Storm services.

### Possible Causes

The heap memory usage is high or the heap memory is improperly allocated.

## Procedure

**Step 1** Check the heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Choose **ALM-26054 Heap Memory Usage of Storm Nimbus Exceeds the Threshold > Location**. Query the **HostName** of the alarmed instance.
3. Choose **Components > Storm > Instances > Nimbus (corresponding to the HostName of the alarmed instance) > Customize > Heap Memory Usage of Nimbus**.
4. Check whether the heap memory usage of Nimbus has reached the threshold (80%).
  - If yes, go to **Step 1.5**.
  - If no, go to **Step 2**.
5. Adjust the heap memory.  
Choose **Components > Storm > Service Configuration**, and set **Type** to **All**. Choose **Nimbus > System**. Increase the value of **-Xmx** in **NIMBUS\_GC\_OPTS**. Click **Save Configuration**. Select **Restart the affected services or instances** and click **OK**.
6. Check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to **Step 2**.

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.499 ALM-27001 DBService Unavailable (For MRS 2.x or Earlier)

### Description

The alarm module checks the DBService status every 30 seconds. This alarm is generated when the system detects that DBService is unavailable.

This alarm is cleared when DBService recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 27001    | Critical       | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

The database service is unavailable and cannot provide data import and query functions for upper-layer services, which results in service exceptions.

## Possible Causes

- The floating IP address does not exist.
- There is no active DBServer instance.
- The active and standby DBServer processes are abnormal.

## Procedure

**Step 1** Check whether the floating IP address exists in the cluster environment.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **DBService > Instances**.
3. Check whether the active instance exists.
  - If yes, go to [Step 1.4](#).
  - If no, go to [Step 2.1](#).
4. Select the active DBServer instance and record the IP address.
5. Log in to the host with the preceding IP address and run the **ifconfig** command to check whether the DBService floating IP address exists on the node.
  - If yes, go to [Step 1.6](#).
  - If no, go to [Step 2.1](#).
6. Run the **ping floating IP address** command to check whether the DBService floating IP address can be pinged.
  - If yes, go to [Step 1.7](#).
  - If no, go to [Step 2.1](#).
7. Log in to the host where the DBService floating IP address is located and run the **ifconfig interface down** command to delete the floating IP address.
8. Choose **Components > DBService > More > Restart Service** to restart DBService and check whether DBService is started successfully.

- If yes, go to [Step 1.9](#).
  - If no, go to [Step 2.1](#).
9. Wait about 2 minutes and check whether the alarm is cleared in the alarm list.
- If yes, no further action is required.
  - If no, go to [Step 13](#).

**Step 2** Check the status of the active DBServer instance.

1. Select the DBServer instance whose role status is abnormal and record the IP address.
2. On the **Alarms** page, check whether ALM-12007 Process Fault occurs in the DBServer instance on the host that corresponds to the IP address.
  - If yes, go to [Step 2.3](#).
  - If no, go to [Step 4](#).
3. Rectify the fault by following steps provided in ALM-12007 Process Fault.
4. Wait about 5 minutes and check whether the alarm is cleared in the alarm list.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 3** Check the status of the active and standby DBServers.

1. Log in to the host where the DBService floating IP address is located, run the **sudo su - root** and **su - omm** commands to switch to user **omm**, and run the **cd \${BIGDATA\_HOME}/FusionInsight/dbservice/** command to go to the DBService installation directory.
2. Run the **sh sbin/status-dbserver.sh** command to view the status of the active and standby HA processes of DBService. Determine whether the status can be viewed successfully.
  - If yes, go to [Step 3.3](#).
  - If no, go to [Step 4](#).
3. Check whether the active and standby HA processes are abnormal.
  - If yes, go to [Step 3.4](#).
  - If no, go to [Step 4](#).
4. Choose **Components > DBService > More > Restart Service** to restart DBService and check whether DBService is started successfully.
  - If yes, go to [Step 3.5](#).
  - If no, go to [Step 4](#).
5. Wait about 2 minutes and check whether the alarm is cleared in the alarm list.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.



2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.500 ALM-27003 DBService Heartbeat Interruption Between the Active and Standby Nodes (For MRS 2.x or Earlier)

### Description

This alarm is generated when the active or standby DBService node does not receive heartbeat messages from the peer node.

This alarm is cleared when the heartbeat recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 27003    | Major          | Yes        |

### Parameters

| Parameter               | Description                                             |
|-------------------------|---------------------------------------------------------|
| ServiceName             | Specifies the service for which the alarm is generated. |
| RoleName                | Specifies the role for which the alarm is generated.    |
| HostName                | Specifies the host for which the alarm is generated.    |
| Local DBService HA Name | Specifies a local DBService HA.                         |
| Peer DBService HA Name  | Specifies a peer DBService HA.                          |

### Impact on the System

During the DBService heartbeat interruption, only one node can provide the service. If this node is faulty, no standby node is available for failover and the service is unavailable.

## Possible Causes

The link between the active and standby DBService nodes is abnormal.

## Procedure

**Step 1** Check whether the network between the active and standby DBService servers is in normal state.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, locate the row that contains the alarm and view the IP address of the standby DBService server in the alarm details.
3. Log in to the active DBService server.
4. Run the **ping heartbeat IP address of the standby DBService** command to check whether the standby DBService server is reachable.
  - If yes, go to **Step 2**.
  - If no, go to **Step 1.5**.
5. Contact the network administrator to check whether the network is faulty.
  - If yes, go to **Step 1.6**.
  - If no, go to **Step 2**.
6. Rectify the network fault and check whether the alarm is cleared from the alarm list.
  - If yes, no further action is required.
  - If no, go to **Step 2**.

**Step 2** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.501 ALM-27004 Data Inconsistency Between Active and Standby DBServices (For MRS 2.x or Earlier)

### Description

The system checks the data synchronization status between the active and standby DBServices every 10 seconds. This alarm is generated when the synchronization status cannot be queried for six consecutive times or when the synchronization status is abnormal.

This alarm is cleared when the synchronization is in normal state.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 27004    | Critical       | Yes        |

## Parameters

| Parameter               | Description                                             |
|-------------------------|---------------------------------------------------------|
| ServiceName             | Specifies the service for which the alarm is generated. |
| RoleName                | Specifies the role for which the alarm is generated.    |
| HostName                | Specifies the host for which the alarm is generated.    |
| Local DBService HA Name | Specifies a local DBService HA.                         |
| Peer DBService HA Name  | Specifies a peer DBService HA.                          |
| SYNC_PERCENT            | Synchronization percentage.                             |

## Impact on the System

When data is not synchronized between the active and standby DBServices, the data may be lost or abnormal if the active instance becomes abnormal.

## Possible Causes

- The network between the active and standby nodes is unstable.
- The standby DBService is abnormal.
- The disk space of the standby node is full.

## Procedure

**Step 1** Check whether the network between the active and standby nodes is in normal state.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, locate the row that contains the alarm and view the IP address of the standby DBService node in the alarm details.
3. Log in to the active DBService node.
4. Run the **ping heartbeat IP address of the standby DBService** command to check whether the standby DBService node is reachable.
  - If yes, go to [Step 2.1](#).
  - If no, go to [Step 1.5](#).

5. Contact the O&M personnel to check whether the network is faulty.
  - If yes, go to [Step 1.6](#).
  - If no, go to [Step 2.1](#).
6. Rectify the network fault and check whether the alarm is cleared from the alarm list.
  - If yes, no further action is required.
  - If no, go to [Step 2.1](#).

**Step 2** Check whether the standby DBService is in normal state.

1. Log in to the standby DBService node.
2. Run the following commands to switch the user:  
**sudo su - root**  
**su - omm**
3. Go to the **\${DBSERVER\_HOME}/sbin** directory and run the **./status-dbserver.sh** command to check whether the GaussDB resource status of the standby DBService is in normal state. In the command output, check whether the following information is displayed in the row where **ResName** is **gaussDB**:  
Example:  

```
10_10_10_231 gaussDB Standby_normal Normal Active_standby
```

  - If yes, go to [Step 3.1](#).
  - If no, go to [Step 4](#).

**Step 3** Check whether the disk space of the standby node is insufficient.

1. Log in to the standby DBService node.
2. Run the following commands to switch the user:  
**sudo su - root**  
**su - omm**
3. Go to the **\${DBSERVER\_HOME}** directory, and run the following commands to obtain the DBService data directory:  
**cd \${DBSERVER\_HOME}**  
**source .dbservice\_profile**  
**echo \${DBSERVICE\_DATA\_DIR}**
4. Run the **df -h** command to check the system disk partition usage.
5. Check whether the DBService data directory space is full.
  - If yes, go to [Step 3.6](#).
  - If no, go to [Step 4](#).
6. Perform upgrade and expand capacity.
7. After capacity expansion, wait 2 minutes and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

## 7.12.502 ALM-28001 Spark Service Unavailable (For MRS 2.x or Earlier)

### Description

The system checks the Spark service status every 30 seconds. This alarm is generated when the Spark service is unavailable.

This alarm is cleared when the Spark service recovers.

### Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 28001    | Critical       | Yes        |

### Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

### Impact on the System

The Spark tasks submitted by users fail to be executed.

### Possible Causes

- The KrbServer service is abnormal.
- The LdapServer service is abnormal.
- The ZooKeeper service is abnormal.
- The HDFS service is abnormal.
- The Yarn service is abnormal.

- The corresponding Hive service is abnormal.

## Procedure

**Step 1** Check whether service unavailability alarms exist in services that Spark depends on.

1. Go to the MRS cluster details page and choose **Alarms**.
2. Check whether the following alarms exist in the alarm list:
  - a. ALM-25500 KrbServer Service Unavailable
  - b. ALM-25000 LdapServer Service Unavailable
  - c. ALM-13000 ZooKeeper Service Unavailable
  - d. ALM-14000 HDFS Service Unavailable
  - e. ALM-18000 Yarn Service Unavailable
  - f. ALM-16004 Hive Service Unavailable
  - If yes, go to [Step 1.3](#).
  - If no, go to [Step 2](#).
3. Handle the alarms based on the troubleshooting methods provided in the alarm help.

After the alarm is cleared, wait a few minutes and check whether the alarm HetuServer Service Unavailable is cleared.

- If yes, no further action is required.
- If no, go to [Step 2](#).

**Step 2** Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Reference

None

### 7.12.503 ALM-38000 Kafka Service Unavailable (For MRS 2.x or Earlier)

#### Description

The system checks the Kafka service availability every 30 seconds. This alarm is generated if the Kafka service becomes unavailable.

This alarm is cleared after the Kafka service recovers.

## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 38000    | Critical       | Yes        |

## Parameters

| Parameter   | Description                                             |
|-------------|---------------------------------------------------------|
| ServiceName | Specifies the service for which the alarm is generated. |
| RoleName    | Specifies the role for which the alarm is generated.    |
| HostName    | Specifies the host for which the alarm is generated.    |

## Impact on the System

The cluster cannot provide the Kafka service and users cannot run new Kafka tasks.

## Possible Causes

- The KrbServer component is faulty.
- The ZooKeeper component is faulty or fails to respond.
- The Broker node in the Kafka cluster is abnormal.

## Procedure

- Step 1** Check the KrbServer component status. For clusters without Kerberos authentication, skip this step and go to [Step 2](#).
1. Go to the MRS cluster details page and click **Components**.
  2. Check whether the health status of the KrbServer service is **Good**.
    - If yes, go to [Step 2.1](#).
    - If no, go to [Step 1.3](#).
  3. Rectify the fault by following instructions in ALM-25500 KrbServer Service Unavailable.
  4. Perform [Step 1.2](#) again.
- Step 2** Check the ZooKeeper component status.
1. Check whether the health status of the ZooKeeper service is **Good**.
    - If yes, go to [Step 3.1](#).
    - If no, go to [Step 2.2](#).

2. If the ZooKeeper service is stopped, start it. For other problems, follow the instructions in ALM-13000 ZooKeeper Service Unavailable.
3. Perform [Step 2.1](#) again.

**Step 3** Check the Broker status.

1. Choose **Components > Kafka > Broker**.
2. In **Role**, check whether all instances are normal.
  - If yes, go to [Step 3.4](#).
  - If no, go to [Step 3.3](#).
3. Select all instances of Broker and choose **More > Restart Instance**.
  - If the restart is successful, go to [Step 3.4](#).
  - If the restart fails, go to [Step 4](#).
4. Choose **Components > Kafka**. Check whether the health status of Kafka is **Good**.
  - If yes, go to [Step 3.5](#).
  - If no, go to [Step 4](#).
5. Wait 30 seconds and check whether the alarm is cleared.
  - If yes, no further action is required.
  - If no, go to [Step 4](#).

**Step 4** Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

## Related Information

N/A

## 7.12.504 ALM-38001 Insufficient Kafka Disk Capacity (For MRS 2.x or Earlier)

### Description

The system checks the Kafka disk usage every 60 seconds and compares it with the threshold. This alarm is generated if the disk usage exceeds the threshold.

To modify the threshold, users can choose **System > Threshold Configuration** on MRS Manager.

This alarm is cleared if the Kafka disk usage is lower than or equal to the threshold.



## Attribute

| Alarm ID | Alarm Severity | Auto Clear |
|----------|----------------|------------|
| 38001    | Major          | Yes        |

## Parameters

| Parameter         | Description                                                                         |
|-------------------|-------------------------------------------------------------------------------------|
| ServiceName       | Specifies the service for which the alarm is generated.                             |
| RoleName          | Specifies the role for which the alarm is generated.                                |
| HostName          | Specifies the host for which the alarm is generated.                                |
| PartitionName     | Specifies the disk partition where the alarm is generated.                          |
| Trigger Condition | Generates an alarm when the actual indicator value exceeds the specified threshold. |

## Impact on the System


Kafka fails to write data to the disks.

## Possible Causes

- The Kafka disk configurations (such as disk count and disk size) are insufficient for the data volume.
- The data retention period is long and historical data occupies large space.
- Services are improperly planned. As a result, data is unevenly distributed and some disks are full.

## Procedure

- Step 1** Go to the MRS cluster details page and choose **Alarms**.
- Step 2** In the alarm list, click the alarm and view the **HostName** and **PartitionName** of the alarm in **Location** of **Alarm Details**.
- Step 3** On the **Hosts** page, click the host name obtained in [Step 2](#).
- Step 4** Check whether the **Disk** area contains the **PartitionName** of the alarm.
  - If yes, go to [Step 5](#).
  - If no, manually clear the alarm and no further action is required.

- Step 5** In the **Disk** area, check whether the usage of the alarmed partition has reached 100%.
- If yes, go to [Step 6](#).
  - If no, go to [Step 8](#).
- Step 6** In **Instance**, choose **Broker > Instance Configuration**. On the **Instance Configuration** page that is displayed, set **Type** to **All** and query the data directory parameter **log.dirs**.
- Step 7** Choose **Components > Kafka > Instances**. On the **Kafka Instance** page that is displayed, stop the Broker instance corresponding to [Step 2](#). Then log in to the alarmed node and manually delete the data directory in [Step 6](#). After all subsequent operations are complete, start the Broker instance.
- Step 8** Choose **Components > Kafka > Service Configuration**. The **Kafka Configuration** page is displayed.
- Step 9** Check whether **disk.adapter.enable** is **true**.
- If yes, go to [Step 11](#).
  - If no, change the value to **true** and go to [Step 10](#).
- Step 10** Check whether the **adapter.topic.min.retention.hours** parameter, indicating the minimum data retention period, is properly configured.
- If yes, go to [Step 12](#).
  - If no, set it to a proper value and go to [Step 12](#).
-  **NOTE**
- If the retention period cannot be adjusted for certain topics, the topics can be added to **disk.adapter.topic.blacklist**.
- Step 11** Wait 10 minutes and check whether the disk usage is reduced.
- If yes, wait until the alarm is cleared.
  - If no, go to [Step 12](#).
- Step 12** Go to the **Kafka Topic Monitor** page and query the data retention period configured for Kafka. Determine whether the retention period needs to be shortened based on service requirements and data volume.
- If yes, go to [Step 13](#).
  - If no, go to [Step 14](#).
- Step 13** Find the topics with great data volumes based on the disk partition obtained in [Step 2](#). Log in to the Kafka client and manually shorten the data retention period for these topics using the following command:
- ```
kafka-topics.sh --zookeeper ZooKeeper address:24002/kafka --alter --topic Topic name --config retention.ms=Retention period
```
- Step 14** Check whether partitions are properly configured for topics. For example, if the number of partitions for a topic with a large data volume is smaller than the number of disks, data may be unevenly distributed to the disks and the usage of some disks will reach the upper limit.

 NOTE

To identify topics with great data volumes, log in to the relevant nodes that are obtained in [Step 2](#), go to the data directory (the directory before `log.dirs` in [Step 6](#) is modified), and check the disk space occupied by the partitions of the topics.

- If the partitions are improperly configured, go to [Step 15](#).
- If the partitions are properly configured, go to [Step 16](#).

Step 15 On the Kafka client, add partitions to the topics.

```
kafka-topics.sh --zookeeper ZooKeeper address:24002/kafka --alter --topic Topic name --partitions=Number of new partitions
```

 NOTE

It is advised to set the number of new partitions to a multiple of the number of Kafka disks. This operation may not quickly clear the alarm. Data will be gradually balanced among the disks.

Step 16 Check whether the cluster capacity needs to be expanded.

- If yes, add nodes to the cluster and go to [Step 17](#).
- If no, go to [Step 17](#).

Step 17 Wait a moment and then check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to [Step 18](#).

Step 18 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

7.12.505 ALM-38002 Heap Memory Usage of Kafka Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the heap memory usage of Kafka every 30 seconds. This alarm is generated if the heap memory usage of Kafka exceeds the threshold (80%).

This alarm is cleared if the heap memory usage is lower than the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
38002	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Memory overflow may occur, causing service crashes.

Possible Causes

The heap memory usage is high or the heap memory is improperly allocated.

Procedure

Step 1 Check the heap memory usage.

1. Go to the MRS cluster details page and choose **Alarms**.
2. Choose **ALM-38002 Kafka Heap Memory Usage Exceeds the Threshold > Location**. Query the IP address of the alarmed instance.
3. Choose **Components > Kafka > Instance > Broker (corresponding to the IP address of the alarmed instance) > Customize > Kafka Heap Memory Resource Percentage** to check the heap memory usage.
4. Check whether the heap memory usage of Kafka has reached the threshold (80%).
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Kafka > Service Configuration > All > Broker > Environment Variables**. Increase the value of **KAFKA_HEAP_OPTS** as required.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

7.12.506 ALM-43001 Spark Service Unavailable (For MRS 2.x or Earlier)

Description

The system checks the Spark service status every 60 seconds. This alarm is generated when the Spark service is unavailable.

This alarm is cleared when the Spark service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43001	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The Spark tasks submitted by users fail to be executed.

Possible Causes

- The KrbServer service is abnormal.
- The LdapServer service is abnormal.
- ZooKeeper is abnormal.
- The HDFS service is abnormal.
- The Yarn service is abnormal.

- The corresponding Hive service is abnormal.

Procedure

Step 1 Check whether service unavailability alarms exist in services that Spark depends on.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether the following alarms exist in the alarm list:
 - a. ALM-25500 KrbServer Service Unavailable
 - b. ALM-25000 LdapServer Service Unavailable
 - c. ALM-13000 ZooKeeper Service Unavailable
 - d. ALM-14000 HDFS Service Unavailable
 - e. ALM-18000 Yarn Service Unavailable
 - f. ALM-16004 Hive Service Unavailable
 - If yes, go to **Step 1.3**.
 - If no, go to **Step 2**.
3. Handle the alarm according to the alarm help.

After the alarm is cleared, wait a few minutes and check whether the alarm HetuServer Service Unavailable is cleared.

 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.507 ALM-43006 Heap Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the JobHistory process status every 30 seconds. The alarm is generated when the heap memory usage of the JobHistory process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43006	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JobHistory process heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the JobHistory process is overused or the heap memory is inappropriately allocated.

Procedure

Step 1 Check the heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43006** and view the IP address and role name of the instance in **Location**.
3. Choose **Components > Spark > Instance > JobHistory** (IP address of the instance for which the alarm is generated) > **Customize > Heap Memory Statistics of the JobHistory Process**. Click **OK** to view the heap memory usage.
4. Check whether the used heap memory of JobHistory reaches 90% of the maximum heap memory specified for JobHistory.
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Spark > Service Configuration**. Set **Type** to **All** and choose **JobHistory > Default**. Increase the value of **SPARK_DAEMON_MEMORY** as required.
6. Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK**.
7. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.508 ALM-43007 Non-Heap Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the JobHistory process status every 30 seconds. The alarm is generated when the non-heap memory usage of the JobHistory process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43007	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JobHistory process non-heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The non-heap memory of the JobHistory process is overused or the non-heap memory is inappropriately allocated.

Procedure

Step 1 Check non-heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43007** and view the IP address and role name of the instance in **Location**.
3. Choose **Components > Spark > Instance > JobHistory** (IP address of the instance for which the alarm is generated) **> Customize > Non-Heap Memory Statistics of the JobHistory Process**. Click **OK** to view the non-heap memory usage.
4. Check whether the non-heap memory usage of JobHistory has reached the threshold (90% of the maximum memory).
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Spark > Service Configuration**. Set **Type** to **All** and choose **JobHistory > Default**. Increase the value of **-XX:MaxMetaspaceSize** in **SPARK_DAEMON_JAVA_OPTS** as required.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.509 ALM-43008 Direct Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the JobHistory process status every 30 seconds. The alarm is generated when the direct memory usage of the JobHistory process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43008	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JobHistory process direct memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the JobHistory process is overused or the direct memory is inappropriately allocated.

Procedure

Step 1 Check the direct memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43008** and view the IP address and role name of the instance in **Location**.
3. Choose **Components > Spark > Instance > JobHistory** (IP address of the instance for which the alarm is generated) **> Customize > Direct Memory Statistics of the JobHistory Process**. Click **OK** to view the direct memory usage.
4. Check whether the direct memory usage of the JobHistory process has reached the threshold (90% of the maximum direct memory).
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Spark > Service Configuration**. Set **Type** to **All** and choose **JobHistory > Default**. Increase the value of - **XX:MaxDirectMemorySize** in **SPARK_DAEMON_JAVA_OPTS** as required.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.510 ALM-43009 JobHistory GC Time Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the GC time of the JobHistory process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (12 seconds) for three consecutive times. You can change the threshold by choosing **System > Threshold Configuration > Service > Spark > JobHistory GC Time > Total JobHistory GC Time**. This alarm is cleared when the JobHistory GC time is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43009	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the GC time exceeds the threshold, JobHistory may run in low performance.

Possible Causes

The heap memory of the JobHistory process is overused or inappropriately allocated, causing frequent GC.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.

2. Select the alarm whose **Alarm ID** is **43009** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Spark** > **Instance** > **JobHistory** (IP address of the instance for which the alarm is generated) > **Customize** > **GC Time of the JobHistory Process**. Click **OK** to view the GC time.
4. Check whether the GC time of the JobHistory process is longer than 12 seconds.
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JobHistory** > **Default**. Increase the value of the **SPARK_DAEMON_MEMORY** parameter as required.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.511 ALM-43010 Heap Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the JDBCServer process status every 30 seconds. The alarm is generated when the heap memory usage of the JDBCServer process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43010	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JDBCServer process heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the JDBCServer process is overused or the heap memory is inappropriately allocated.

Procedure

Step 1 Check the heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43010** and view the IP address and role name of the instance in **Location**.
3. Choose **Components > Spark > Instance > JDBCServer** (IP address of the instance for which the alarm is generated) **> Customize > Heap Memory Statistics of the JDBCServer Process**. Click **OK** to view the heap memory usage.
4. Check whether the heap memory usage of JDBCServer has reached the threshold (90% of the maximum heap memory).
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Spark > Service Configuration**. Set **Type** to **All** and choose **JDBCServer > Tuning**. Increase the value of the **SPARK_DRIVER_MEMORY** parameter as required.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.512 ALM-43011 Non-Heap Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the JDBCServer process status every 30 seconds. The alarm is generated when the non-heap memory usage of the JDBCServer process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43011	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JDBCServer process non-heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The non-heap memory of the JDBCServer process is overused or the non-heap memory is inappropriately allocated.

Procedure

Step 1 Check non-heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43011** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Spark** > **Instance** > **JDBCServer** (IP address of the instance for which the alarm is generated) > **Customize** > **Non-heap**

Memory Statistics of the JDBCServer Process. Click **OK** to view the non-heap memory usage.

4. Check whether the non-heap memory usage of JDBCServer has reached the threshold (90% of the maximum non-heap memory).
 - If yes, go to [Step 1.5](#).
 - If no, go to [Step 2](#).
5. Choose **Components > Spark > Service Configuration**. Set **Type** to **All** and choose **JDBCServer > Tuning**. Increase the value of **-XX:MaxMetaspaceSize** in **spark.driver.extraJavaOptions** as required.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to [Step 2](#).

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.513 ALM-43012 Direct Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the JDBCServer process status every 30 seconds. The alarm is generated when the direct memory usage of the JDBCServer process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43012	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JDBCServer process direct memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the JDBCServer process is overused or the direct memory is inappropriately allocated.

Procedure

Step 1 Check the direct memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43012** and view the IP address and role name of the instance in **Location**.
3. Choose **Components > Spark > Instance > JDBCServer** (IP address of the instance for which the alarm is generated) **> Customize > Direct Memory Statistics of the JDBCServer Process**. Click **OK** to view the direct memory usage.
4. Check whether the direct memory usage of the JDBCServer process has reached the threshold (90% of the maximum direct memory).
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Spark > Service Configuration**. Set **Type** to **All** and choose **JDBCServer > Tuning**. Increase the value of - **XX:MaxDirectMemorySize** in **spark.driver.extraJavaOptions** as required.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.514 ALM-43013 JDBCServer GC Time Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system checks the GC time of the JDBCServer process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (12 seconds) for three consecutive times. You can change the threshold by choosing **System > Threshold Configuration > Service > Spark > JDBCServer GC Time > Total JDBCServer GC Time**. This alarm is cleared when the JDBCServer GC time is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43013	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the GC time exceeds the threshold, JDBCServer may run in low performance.

Possible Causes

The heap memory of the JDBCServer process is overused or inappropriately allocated, causing frequent GC.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43013** and view the IP address and role name of the instance in **Location**.

3. Choose **Components > Spark > Instance > JDBCServer** (IP address of the instance for which the alarm is generated) > **Customize > GC Time of the JDBCServer Process**. Click **OK** to view the GC time.
4. Check whether the GC time of the JDBCServer process is longer than 12 seconds.
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Spark > Service Configuration**. Set **Type** to **All** and choose **JDBCServer > Tuning**. Increase the value of the **SPARK_DRIVER_MEMORY** parameter as required.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.515 ALM-44004 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold (For MRS 2.x or Earlier)

Description

This alarm is generated when the system detects that the number of queuing tasks in a resource group exceeds the threshold. The system queries the number of queuing tasks in a resource group through the JMX interface. You can choose **Components > Presto > Service Configuration** (switch **Basic** to **All**) > **Presto > resource-groups** to configure a resource group. You can choose **Components > Presto > Service Configuration** (switch **Basic** to **All**) > **Coordinator > Customize > resourceGroupAlarm** to configure the threshold of each resource group.

Attribute

Alarm ID	Alarm Severity	Auto Clear
44004	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

If the number of queuing tasks in a resource group exceeds the threshold, a large number of tasks may be in the queuing state. The Presto task time exceeds the expected value. When the number of queuing tasks in a resource group exceeds the maximum number (**maxQueued**) of queuing tasks in the resource group, new tasks cannot be executed.

Possible Causes

The resource group configuration is improper or too many tasks in the resource group are submitted.

Procedure

- Step 1** Choose **Components > Presto > Service Configuration** (switch **Basic** to **All**) > **Presto > resource-groups** to adjust the resource group configuration.
- Step 2** You can choose **Components > Presto > Service Configuration** (switch **Basic** to **All**) > **Coordinator > Customize > resourceGroupAlarm** to modify the threshold of each resource group.
- Step 3** Collect fault information.
1. Log in to the cluster node based on the host name in the fault information and query the number of queuing tasks based on **Resource Group** in the additional information on the Presto client.
 2. Log in to the cluster node based on the host name in the fault information, view the **/var/log/Bigdata/nodeagent/monitorlog/monitor.log** file, and search for resource group information to view the monitoring collection information of the resource group.
 3. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.516 ALM-44005 Presto Coordinator Process GC Time Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system collects GC time of the Presto Coordinator process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). You can change the threshold by choosing **System > Configure Alarm Threshold > Service > Presto > Coordinator > Presto Process Garbage Collection Time > Garbage Collection Time of the Coordinator Process** on MRS Manager. This alarm is cleared when the Coordinator process GC time is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
44005	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

If the GC time of the Coordinator process is too long, the Coordinator process running performance will be affected and the Coordinator process will even be unavailable.

Possible Causes

The heap memory of the Coordinator process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **44005** and view the IP address and role name of the instance in **Location**.

3. Choose **Components > Presto > Instances > Coordinator** (business IP address of the instance for which the alarm is generated) > **Customize > Presto Garbage Collection Time**. Click **OK** to view the GC time.
4. Check whether the GC time of the Coordinator process is longer than 5 seconds.
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Presto > Service Configuration**, and switch **Basic** to **All**. Choose **Presto > Coordinator**. Increase the value of **-Xmx** (maximum heap memory) in the **JAVA_OPTS** parameter based on the site requirements.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.517 ALM-44006 Presto Worker Process GC Time Exceeds the Threshold (For MRS 2.x or Earlier)

Description

The system collects GC time of the Presto Worker process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). You can change the threshold by choosing **System > Configure Alarm Threshold > Service > Presto > Worker > Presto Garbage Collection Time > Garbage Collection Time of the Worker Process** on MRS Manager. This alarm is cleared when the Worker process GC time is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
44006	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

If the GC time of the Worker process is too long, the Worker process running performance will be affected and the Worker process will even be unavailable.

Possible Causes

The heap memory of the Worker process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **44006**. Then check the IP address and role name of the instance in **Location**.
3. Choose **Components > Presto > Instances > Worker** (business IP address of the instance for which the alarm is generated) **> Customize > Presto Garbage Collection Time**. Click **OK** to view the GC time.
4. Check whether the GC time of the Worker process is longer than 5 seconds.
 - If yes, go to **Step 1.5**.
 - If no, go to **Step 2**.
5. Choose **Components > Presto > Service Configuration**, and switch **Basic** to **All**, and choose **Presto > Worker** Increase the value of **-Xmx** (maximum heap memory) in the **JAVA_OPTS** parameter based on the site requirements.
6. Check whether the alarm is cleared.
 - If yes, no further action is required.
 - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System > Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

7.12.518 ALM-45325 Presto Service Unavailable (For MRS 2.x or Earlier)

Description

The system checks the Presto service status every 60 seconds. This alarm is generated when the system detects that Presto is unavailable.

This alarm is cleared when the Presto service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45325	Critical	Yes

Parameters

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Presto cannot run SQL queries.

Possible Causes

- The Presto coordinator or worker process is faulty.
- The network communication between Presto coordinator and worker instances is interrupted.


Procedure

Step 1 Check the status of the coordinator and worker processes.

1. Log in to FusionInsight Manager and choose **Cluster > Services > Presto**. On the page that is displayed, click the **Instance** tab. In the Presto instance list, check whether the status of all coordinator or worker instances is **Unknown**.
 - If yes, go to [2](#).
 - If no, go to [1](#).

2. In the upper part of the Presto instance list, choose **More > Restart Service** to restart the coordinator and worker processes.
3. In the alarm list, check whether ALM-45325 Presto Service Unavailable is cleared.
 - If yes, no further action is required.
 - If no, go to **1** in **Step 2**.

Step 2 Collect fault information.

1. On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log > Download**.
2. Select **Presto** for **Service**.
3. Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
4. Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

7.13 Configuring Remote O&M for an MRS Cluster

If you encounter any issues while using a cluster and require assistance from Huawei Cloud support personnel, you can first contact them. Then, you can authorize them to access your machine through the O&M authorization function to locate the fault. Alternatively, you can provide Huawei Cloud support personnel with logs from a specific time period using the log sharing function to help them locate the fault.

Enabling Remote O&M Authorization for an MRS Cluster

Step 1 Log in to the MRS management console.

Step 2 In the navigation pane of the MRS management console, choose **Active Clusters**, select a running cluster, and click its name to switch to the cluster details page.

Step 3 In the upper right corner of the page, click **O&M**, choose **Authorize for Cluster Nodes**, and select the deadline for the HUAWEI CLOUD support engineers to access the local host. Before the deadline, the support personnel have the temporary permission to access the local host.

Step 4 Select the check box to confirm authorization and click **OK**.

Step 5 After the fault is rectified, click **O&M** in the upper right corner of the page and choose **Deauthorize for Cluster Nodes** to retrieve the access permission granted to the HUAWEI CLOUD support engineers.

----End

Configuring Remote O&M Log Sharing for an MRS Cluster

- Step 1** Log in to the MRS management console.
- Step 2** In the navigation pane of the MRS management console, choose **Active Clusters**, select a running cluster, and click its name to switch to the cluster details page.
- Step 3** In the upper right corner of the displayed page, choose **O&M > Share Log**.
- Step 4** In the displayed dialog box, select the start time and end time in **Time Range**.

NOTE

- Select **Time Range** based on the suggestions of Huawei Cloud support personnel.
- To filter logs by time, make sure to set **End Date** to a date that comes after **Start Date**.

----End

7.14 Common Ports for MRS Cluster Services

When you [buy a custom cluster](#) of an LTS version, you can customize the component port. If you do not want to customize a port, an open source port is used.

- **Open source:** Find the default port of the component in the Default Open Source Port column of the following table.
- **Custom:** Find the default port of the component in the Default Custom Port column of the following table.
- If there is only the Default Port column, the open source port of the component is the same as the default custom port.

If the cluster is not of an LTS version, the **Component Port** parameter is unavailable and only an open source port can be used. For details, see the Default Open Source Port or Default Port column.

Common HBase Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
hbase.master.port	16000	21300	<p>HMaster RPC port. This port is used to connect the HBase client to HMaster.</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hbase.master.info.port	16010	21301	<p>HMaster HTTPS port. This port is used by the remote web client to connect to the HMaster UI.</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hbase.regionserver.port	16020	21302	<p>RegionServer (RS) RPC port. This port is used to connect the HBase client to RegionServer.</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
hbase.regionserver.info.port	16030	21303	<p>HTTPS port of the Region server. This port is used by the remote web client to connect to the RegionServer UI.</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hbase.thrift.info.port	9095	21304	<p>Thrift Server listening port of Thrift Server</p> <p>This port is used for: Listening when the client is connected</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hbase.regionserver.thrift.port	9090	21305	<p>Thrift Server listening port of RegionServer</p> <p>This port is used for: Listening when the client is connected to the RegionServer</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
hbase.rest.info.port	8085	21308	Port of the RegionServer RESTServer native web page
-	21309	21309	REST port of RegionServer RESTServer

Common HDFS Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
dfs.namenode.rpc.port	<ul style="list-style-type: none"> 9820 (versions earlier than MRS 3.x) 8020 (MRS 3.x or later) 	25000	<p>NameNode RPC port</p> <p>This port is used for:</p> <ol style="list-style-type: none"> 1. Communication between the HDFS client and NameNode 2. Connection between the DataNode and NameNode <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.namenode.http.port	9870	25002	<p>HDFS HTTP port (NameNode)</p> <p>This port is used for:</p> <ol style="list-style-type: none"> 1. Point-to-point NameNode checkpoint operations 2. Connecting the remote web client to the NameNode UI <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
dfs.namenode.https.port	9871	25003	<p>HDFS HTTPS port (NameNode)</p> <p>This port is used for:</p> <ol style="list-style-type: none"> 1. Point-to-point NameNode checkpoint operations 2. Connecting the remote web client to the NameNode UI <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.datanode.ipc.port	9867	25008	<p>IPC server port of DataNode</p> <p>This port is used for:</p> <p>Connection between the client and DataNode to perform RPC operations.</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.datanode.port	9866	25009	<p>DataNode data transmission port</p> <p>This port is used for:</p> <ol style="list-style-type: none"> 1. Transmitting data from HDFS client from or to the DataNode 2. Point-to-point DataNode data transmission <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
dfs.datanode.http.port	9864	25010	<p>DataNode HTTP port</p> <p>This port is used for:</p> <p>Connecting to the DataNode from the remote web client in security mode</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.datanode.https.port	9865	25011	<p>HTTPS port of DataNode</p> <p>This port is used for:</p> <p>Connecting to the DataNode from the remote web client in security mode</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.JournalNode.rpc.port	8485	25012	<p>RPC port of JournalNode</p> <p>This port is used for:</p> <p>Client communication to access multiple types of information</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
dfs.journalnode.http.port	8480	25013	<p>JournalNode HTTP port</p> <p>This port is used for:</p> <p>Connecting to the JournalNode from the remote web client in security mode</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.journalnode.https.port	8481	25014	<p>HTTPS port of JournalNode</p> <p>This port is used for:</p> <p>Connecting to the JournalNode from the remote web client in security mode</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
httpfs.http.port	14000	25018	<p>Listening port of the HttpFS HTTP server</p> <p>This port is used for:</p> <p>Connecting to the HttpFS from the remote REST API</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common HetuEngine Ports

The protocol type of the port in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
server.port (HSBroker)	29860	29860	Specifies the port number that HSBroker listens to.
server.port (HSConsole)	29880	29880	Specifies the port number that HSConsole listens to.
server.port (HSFabric)	29900	29900	Specifies the port number that HSFabric listens to, which is used for cross-domain connections
gateway.port	29902	29902	Specifies the port number that HSFabric listens to, which is used for JDBC connections

Common Hive Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
templeton.port	9111	21055	<p>Port used for WebHCat to provide the REST service</p> <p>This port is used for:</p> <p>Communication between the WebHCat client and WebHCat server</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
hive.server2.thrift.port	10000	21066	<p>Port for HiveServer to provide Thrift services</p> <p>This port is used for:</p> <p>Communication between the HiveServer and HiveServer client</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hive.metastore.port	9083	21088	<p>Port for MetaStore to provide Thrift services</p> <p>This port is used for:</p> <p>Communication between the MetaStore client and MetaStore, that is, communication between HiveServer and MetaStore.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hive.server2.webui.port	10002	-	<p>Web UI port of Hive</p> <p>This port is used for HTTPS/HTTP communication between web requests and the Hive UI server.</p>

Common Hue Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
HTTP_PORT	8888	21200	<p>Port for Hue to provide HTTPS services</p> <p>This port is used to provide web services in HTTPS mode, which can be changed.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common Kafka Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
port	9092	21005	Port for a broker to receive data and obtain services
ssl.port	9093	21008	SSL port used by a broker to receive data and obtain services
sasl.port	21007	21007	SASL security authentication port provided by a broker, which provides the secure Kafka service
sasl-ssl.port	21009	21009	Port used by a broker to provide encrypted service based on the SASL and SSL protocols

Common Loader Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Port	Port Description
LOADER_HTTPS_PORT	21351	<p>This port is used to provide REST APIs for configuration and running of Loader jobs.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common Manager Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Port (Versions Earlier Than MRS 3.x)	Port Description
-	8080	<p>Port provided by WebService for user access</p> <p>This port is used to access the web UI over HTTP.</p> <ul style="list-style-type: none"> Is the port enabled by default during the installation: Yes Is the port enabled after security hardening: Yes
-	28443	<p>Port provided by WebService for user access</p> <p>This port is used to access the web UI over HTTPS.</p> <ul style="list-style-type: none"> Is the port enabled by default during the installation: Yes Is the port enabled after security hardening: Yes

Common MapReduce Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
mapreduce.jobhistory.webapp.port	19888	26012	<p>Web HTTP port of the JobHistory server</p> <p>This port is used for: viewing the web page of the JobHistory server</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> Is the port enabled by default during the installation: Yes Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
mapreduce.jobhistory.port	10020	26013	<p>Port of the JobHistory server</p> <p>This port is used for:</p> <ol style="list-style-type: none"> 1. Task data restoration in the MapReduce client 2. Obtaining task report in the Job client <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
mapreduce.jobhistory.webapp.https.port	19890	26014	<p>Web HTTPS port of the JobHistory server</p> <p>This port is used to view the web page of the JobHistory server.</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common Spark Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
hive.server2.thrift.port	22550	22550	<p>JDBC thrift port</p> <p>This port is used for: Socket communication between Spark 2.1.0 CLI/JDBC client and server</p> <p>NOTE If hive.server2.thrift.port is occupied, an exception indicating that the port is occupied is reported.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
spark.ui.port	4040	22950	<p>Web UI port of JDBC</p> <p>This port is used for: HTTPS/HTTP communication between Web requests and the JDBC Server Web UI server</p> <p>NOTE The system verifies the port configuration. If the port is invalid, the value of the port plus 1 is used till the calculated value is valid. (A maximum number of 16 attempts are allowed. The number of attempts is specified by spark.port.maxRetries.)</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
spark.history.ui.port	18080	22500	<p>JobHistory Web UI port</p> <p>This port is used for: HTTPS/HTTP communication between Web requests and Spark2.1.0 History Server</p> <p>NOTE The system verifies the port configuration. If the port is invalid, the value of the port plus 1 is used till the calculated value is valid. (A maximum number of 16 attempts are allowed. The number of attempts is specified by spark.port.maxRetries.)</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common Storm Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
nimbus.thrift.port	6627	29200	Port for Nimbus to provide thrift services
supervisor.slots.ports	6700,6701,6702,6703	29200-29499	Port for receiving service requests that are forwarded from other servers
logviewer.https.port	29248	29248	Port for LogViewer to provide HTTPS services
ui.https.port	29243	29243	Port for Storm UI to provide HTTPS services (ui.https.port)

Common YARN Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
yarn.resourcemanager.webapp.port	8088	26000	Web HTTP port of the ResourceManager service
yarn.resourcemanager.webapp.https.port	8090	26001	<p>Web HTTPS port of the ResourceManager service</p> <p>This port is used to access the Resource Manager web applications in security mode.</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
yarn.nodemanager.webapp.port	8042	26006	NodeManager Web HTTP port

Parameter	Default Open Source Port	Default Custom Port	Port Description
yarn.nodemanager.webapp.https.port	8044	26010	<p>NodeManager Web HTTPS port</p> <p>This port is used for:</p> <p>Accessing the NodeManager web application in security mode</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common ZooKeeper Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
clientPort	2181	24002	<p>ZooKeeper client port</p> <p>This port is used for:</p> <p>Connection between the ZooKeeper client and server.</p> <p>NOTE</p> <p>The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common Kerberos Ports

The protocol type of all ports in the table is TCP and UDP.

Parameter	Default Port	Port Description
KADMIN_PORT	21730	<p>Kerberos user management port</p> <p>This port is used for user management.</p> <ul style="list-style-type: none"> • Whether the port is enabled by default during the installation: Yes • Whether the port is enabled after security hardening: Yes
KPASSWD_PORT	21731	<p>Kerberos password changing port</p> <p>This port is used for user management.</p> <ul style="list-style-type: none"> • Whether the port is enabled by default during the installation: Yes • Whether the port is enabled after security hardening: Yes
kdc_ports	21732	<p>Kerberos server port</p> <p>This port is used for performing Kerberos authentication for components.</p> <p>This parameter may be used during the configuration of mutual trust between clusters.</p> <p>NOTE The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code.</p> <ul style="list-style-type: none"> • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common OpenTSDB Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
tsd.network.port	4242	<p>Web UI port of OpenTSDB</p> <p>This port is used for HTTPS/HTTP communication between web requests and the OpenTSDB UI server.</p>

Common Tez Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
tez.ui.port	28888	Web UI port of Tez

Common KafkaManager Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
kafka_manager_port	9099	Web UI port of KafkaManager

Common Presto Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
http-server.http.port	7520	HTTP port for Presto coordinator to provide services to external systems
http-server.https.port	7521	HTTPS port for Presto coordinator to provide services to external systems
http-server.http.port	7530	HTTP port for Presto worker to provide services to external systems
http-server.https.port	7531	HTTPS port for Presto worker to provide services to external systems

Common Flink Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
jobmanager.web.port	32261-32325	Web UI port of Flink This port is used for: HTTP/HTTPS communication between the client web requests and Flink server

Common ClickHouse Ports

The protocol type of the port in the table is TCP and HTTP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
interserver_http_port	9009	9009	HTTP port for the communication between ClickHouse servers.
interserver_https_port	9010	9010	HTTPS port for the communication between ClickHouse servers.
http_port	8123	8123	Port for connecting to the ClickHouse server through HTTP.
https_port	8443	8443	Port for connecting to the ClickHouse server through HTTPS.
tcp_port	9000	9000	Port for connecting the client to the ClickHouse server through TCP.
tcp_port_secure	9440	9440	Port for connecting the client to the ClickHouse server through TCP SSL.
lb_tcp_port	21424	21424	TCP communication port number for ClickHouseBalancer.
lb_http_port	21425	21425	HTTP communication port number for ClickHouseBalancer.
lb_https_port	21426	21426	HTTPS communication port number for ClickHouseBalancer.
lb_tcp_secure_port	21428	21428	TCP SSL communication port number for ClickHouseBalancer.

Common Impala Ports

The protocol type of the port in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
--beeswax_port	21000	21970	Port for impala-shell communication
--hs2_port	21050	21971	Port for Impala application communication
--hs2_http_port	28000	21981	Port used by Impala to provide the HiveServer2 protocol for external systems

Common Doris Ports

The protocol type of the port in the table is TCP and HTTP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
http_port	8030	29980	HTTP port of the FE service
https_port	8050	29991	HTTPS port of the FE service
query_port	9030	29982	Port used by the Doris FE to query connections through the MySQL protocol
rpc_port	9020	29981	Thrift Server port of the FE service
be_port	9060	29984	Thrift Server Port on BE for receiving requests from FE
brpc_port	8060	29987	BRPC port on BE, which is used for communications between BE instances.
heartbeat_service_port	9050	29985	Thrift heartbeat service port on BE, which is used to receive heartbeat messages from FE
webserver_port	8040	29986	HTTP server port on BE
broker_ipc_port	8000	29990	Thrift Server communication port on Broker, which is used to receive requests.

Parameter	Default Open Source Port	Default Custom Port	Port Description
single_replica_load_brpc_port	9070	29988	RPC port used for the communication between the master and slave replicas to import single-replica data
single_replica_load_download_port	8050	29989	Port used by the slave replica to download data files from the master replica through HTTP for single-copy data import

8 Configuring Storage-Compute Decoupling for an MRS Cluster

8.1 Configuration Process

With MRS, you can store data in OBS and dedicate MRS clusters solely to computing tasks, isolating storage and compute resources. This approach offers flexible, on-demand scaling at a lower cost, making it well-suited for big data processing.

NOTE

- In big data storage-computing decoupling scenarios, make sure to use an OBS parallel file system. For details, see [Parallel File System](#). Using a regular object bucket can significantly impact the performance of the cluster.
- If a cluster has been connected to OBS (storage and compute decoupling or cold and hot data separation), you need to manually delete service data on OBS after deleting a component or MRS cluster.

Perform the following steps to use the storage-compute decoupling function:

1. Configure a cluster with decoupled storage and compute.

Select one of the following configurations (Using an agency is recommended.):

- Bind an agency of the ECS type to an MRS cluster to access OBS, preventing the AK/SK from being exposed in the configuration file. For details, see [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#).
- Configure the AK/SK in an MRS cluster. The AK/SK will be exposed in the configuration file in plaintext. Exercise caution when performing this operation. For details, see [How Do I Connect an MRS Cluster Client to OBS Using an AK/SK Pair?](#).

2. Use the cluster.

After the required permissions for accessing OBS are obtained, components in the MRS cluster can access the corresponding files through the client.

For details about how to configure components to access OBS, see the following content:

- [Example for Interconnecting a Cluster Service with OBS](#)

Configuring Storage-Compute Decoupling

1. Create an MRS cluster.

The MRS cluster must contain basic components such as Guardian, Ranger, and Hadoop.

NOTE

Currently, only MRS 3.3.0-LTS or later supports interconnection with OBS using Guardian.

2. Create an OBS agency.

Create an agency with OBS access permissions, which allows Guardian to connect to OBS.

3. Enable Guardian to connect to OBS and set parameters.

Modify Guardian configuration parameters and configure IAM agency authentication information.

4. Configure the policy for clearing component data in the recycle bin directory.

In storage and computing decoupling scenarios, the data anti-deletion function is enabled for components that connect to OBS by default. When you delete data, the deleted object will be moved to the corresponding recycle bin directory. To avoid the risk of storage space being used up, you need to configure a lifecycle policy for the corresponding directory in the OBS file system.

5. Interconnect components with OBS.

6. Components in the MRS cluster can directly access the corresponding path after being granted required permissions for accessing OBS buckets. You can directly access resources in the OBS file system via the component client using an absolute path.

Granting OBS Permissions

Enabling storage-compute decoupling and Ranger authentication for an MRS cluster that connects to OBS using Guardian allows Ranger administrators to grant cluster users permissions to read and write OBS directories and files.

In addition, based on the Guardian permission model, storage-compute decoupling, and Hive cascading authorization, users can be granted service table authorization based on Ranger, which automatically associates fine-grained OBS storage directory permissions without requiring secondary authorization. Users only need to authorize the service table once on Ranger, and the system will automatically associate the permissions of the data storage source in a fine-grained manner, without the need to be aware of the storage path of the table or perform secondary authorization.

 NOTE

- On Ranger, only users in Ranger custom user groups can be granted OBS permissions. Enter 1 to 52 characters for each user group name. Only letters, numbers, underscores (_), and number signs (#) are allowed. Or, the policy fails to add.
- For clusters with Kerberos authentication enabled, permissions need to be granted based on Ranger. For clusters with Kerberos authentication disabled, OBS permissions are granted by default, and no additional configuration is required.
- If Kerberos authentication is not enabled for the current cluster, the user who accesses OBS must belong to the **supergroup** group.

8.2 Interconnecting an MRS Cluster with OBS Using an IAM Agency

8.2.1 Interconnecting an MRS Cluster with OBS Using an IAM Agency

MRS allows you to store data in OBS and use an MRS cluster for data computing only. In this way, storage and compute are separated. You can create an IAM agency, which enables ECS to automatically obtain the temporary AK/SK to access OBS. This prevents the AK/SK from being exposed in the configuration file.

By binding an agency, ECSs or BMSs can manage some of your resources. Determine whether to configure an agency based on the actual service scenario. This feature can be used with Hadoop, Hive, Spark, Presto, and Flink components in clusters. To interconnect MRS with OBS using an IAM agency, perform the following tasks:

1. [Creating an ECS Agency with OBS Access Permissions](#)
2. [Creating a Decoupled Storage and Compute Cluster](#)
3. [Creating an OBS File System for Storing Data](#)
4. [Creating a Lifecycle Rule](#)

Creating an ECS Agency with OBS Access Permissions

 NOTE

- MRS presets **MRS_ECS_DEFAULT_AGENCY** in the IAM agency list by default, allowing you to choose this agency when creating a cluster. This agency has **OBSOperateAccess** permission and, for users with fine-grained policies enabled, **CESFullAccess**, **CES Administrator**, and **KMS Administrator** permissions in the region where the cluster is located. Do not modify **MRS_ECS_DEFAULT_AGENCY** on IAM.
- If you want to use the preset agency, skip the step for creating an agency. If you want to use a custom agency, perform the following steps to create an agency. (To create or modify an agency, you must have the Security Administrator permission.) If you need to have more fine-grained control over the permissions of a specific path in the OBS file system, you can refer to [Configuring Fine-Grained OBS Access Permissions for MRS Cluster Users](#) to create a custom role policy.

1. Log in to the Huawei Cloud management console.
2. In the service list, choose **Management & Governance > Identity and Access Management**.

3. Choose **Agencies**. On the displayed page, click **Create Agency**.
4. Set **Agency Name**. For example, enter `mrs_ecs_obs`.
5. Set **Agency Type** to **Cloud service** and select **ECS BMS** to authorize ECS or BMS to invoke OBS. See [Figure 8-1](#).
6. Set **Validity Period** to **Unlimited** and click **Done**.

Figure 8-1 Creating an agency

* Agency Name

* Agency Type Account
Delegate another Huawei Cloud account to perform operations on your resources.
 Cloud service
Delegate a cloud service to access your resources in other cloud services.

* Cloud Service

* Validity Period

Description
0/255

7. In the displayed dialog box, click **Authorize**. Search for **OBS OperateAccess** and select it.

NOTE

If KMS encryption is configured for an OBS bucket, the **KMS Administrator** policy must be selected.

Figure 8-2 Configuring permissions

Select Policy/Role | Select Scope | Finish

Assign selected permissions to Create Policy

View Selected (1) | Copy Permissions from Another Project | All policies/roles | All services | Fuzzy se... | OBS OperateAccess | X | Q

Policy/Role Name	Type
<input checked="" type="checkbox"/> OBS OperateAccess Basic operation permissions to view the bucket list, obtain bucket metadata, list objects in a bucket, query bucket location, upload objects, download objects, delete...	System-defined policy

8. Click **Next**. On the page that is displayed, select the desired scope for the permissions you selected. By default, **All resources** is selected. Click **Show More**, select **Global resources**, and click **OK**.
9. In the dialog box that is displayed, click **OK** to start authorization. After the message "**Authorization successful.**" is displayed, click **Finish**. The agency is successfully created.

Creating a Decoupled Storage and Compute Cluster

You can configure an agency when creating a cluster or bind an agency to an existing cluster to separate storage and compute. This section uses a cluster with Kerberos authentication enabled as an example.

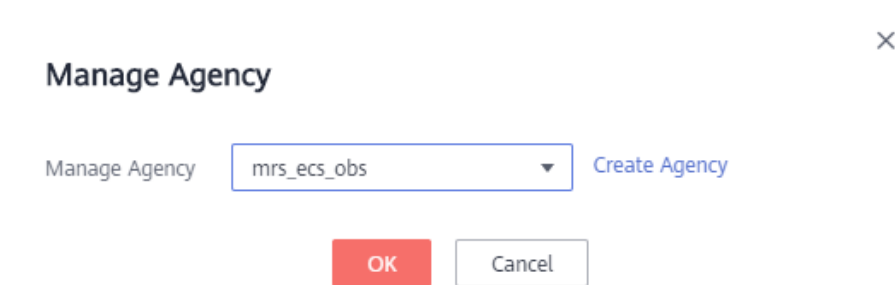
Configuring an agency when creating a cluster:

1. Go to the [Buy Cluster](#) page.
2. Click **Buy Cluster**. The page for buying a cluster is displayed.
3. Click the **Custom Config** tab.
4. On the **Quick Config** tab page, set the following parameters:
 - Basic configuration:
 - **Billing Mode:** Select **Pay-per-use**.
 - **Region:** Select a region as required.
 - Cluster configuration:
 - **Cluster Name:** You can use the default name. However, you are advised to include an abbreviation of the project name or date to make it easier to distinguish and consolidate memory.
 - **Cluster Type:** Select **Custom**.
 - **Version Type:** Select **LTS** or **Normal**.
 - **Cluster Version:** Select a cluster version as needed, for example, **MRS 3.2.0-LTS.1**.
 - **Component:** Be careful when selecting a cluster type that combines multiple components, as certain cluster types do not allow for the addition of components after the cluster has been created.
 - **Metadata:** Select **Local**.
 - Network configuration:
 - **AZ:** Retain the default value.
 - **VPC:** Use the default value.
 - **Subnet:** Use the default value.
 - **Security Group:** Use the default value.
 - **EIP:** Retain the default value.
 - Node configuration:
 - **CPU Architecture:** Retain the default value. This parameter is not available for MRS 3.1.0 and 3.1.5.
 - **Common Template:** This parameter is available only when **Cluster Type** is set to **Custom**. Retain the default value.

- **Cluster Node:** Select the number of cluster nodes and node specifications based on site requirements.
 - Login credentials:
 - **Kerberos Authentication:** Determine whether to enable it as needed. If the cluster to create contains Presto, Kerberos authentication cannot be enabled.
 - **Username:** The default username is **admin**, which is used to log in to FusionInsight Manager.
 - **Password/Confirm Password:** Set a password for the user **admin**. Keep the password secure.
 - **Login Mode:** Select a method for logging in to ECSs. In this example, select **Password**.
 - **Username:** The default username is **root**, which is used to remotely log in to ECSs.
 - **Password/Confirm Password:** Set the password for the user **root**.
 - **Advanced Configuration:** Enable advanced settings and set an agency.
 - Click **Available agencies** and select the agency created in [Creating an ECS Agency with OBS Access Permissions](#) from the drop-down list.
 - Select the **MRS_ECS_DEFAULT_AGENCY** agency preset by MRS in IAM.
 - **Enterprise Project:** Retain the default value.
 - **Secure Communications:** Select this option. For details, see [Configuring Secure Communication Authorization for an MRS Cluster](#).
5. Click **Buy Now** and wait until the cluster is created.
- If Kerberos authentication is enabled for a cluster, check whether Kerberos authentication is required. If yes, click **Continue**. If no, click **Back** to disable Kerberos authentication and then create a cluster.

Configuring an agency for an existing cluster:

1. Log in to the MRS console. In the navigation pane on the left, choose **Active Clusters**.
2. Click the name of the cluster to enter its details page.
3. On the **Dashboard** page, click **Synchronize** on the right of **IAM User Sync** to synchronize IAM users.
4. On the **Dashboard** tab page, click **Manage Agency** on the right side of **Agency** to select an agency and click **OK** to bind it. Alternatively, click **Create Agency** to go to the IAM console to create an agency and select it.

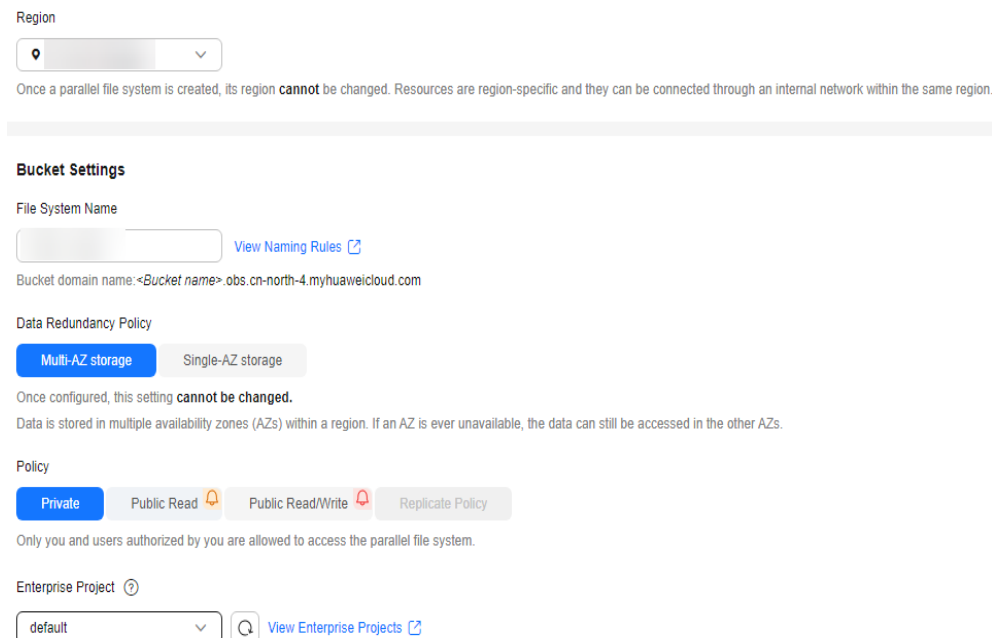
Figure 8-3 Binding an agency

Creating an OBS File System for Storing Data

NOTE

In big data storage-computing decoupling scenarios, make sure to use an OBS parallel file system. For details, see [Parallel File System](#). Using a regular object bucket can significantly impact the performance of the cluster.

1. Log in to the OBS Console.
2. Choose **Parallel File Systems** > **Create Parallel File System**.
3. Enter the file system name, for example, **mrs-word001**.
Set other parameters as required.

Figure 8-4 Creating an OBS parallel file system

4. Click **Create Now**.
5. In the parallel file system list on the OBS console, click the file system name to go to the details page.
6. In the navigation pane, choose **Files** and create the **program** and **input** folders.

- **program**: Upload the program package to this folder.
- **input**: Upload the input data to this folder.

Creating a Lifecycle Rule

In MRS 3.2.0-LTS.1 and later versions, components prevent mis-deletion by default. That is, file data deleted by component users is not directly deleted but stored in the recycle bin directory in the OBS file system.

To save OBS space, you need to enable periodical deletion of file data from the OBS recycle bin by referring to [Configuring the Policy for Clearing Recycle Bin Directories of MRS Cluster Components](#).

8.2.2 Configuring the Policy for Clearing Recycle Bin Directories of MRS Cluster Components

Scenario

By default, components in an MRS 3.2.0-LTS.1 or later cluster support prevention against accidental data deletion. Native HDFS garbage collection can be used in the Hadoop big data systems that use OBS.

The file data deleted by a component user is not directly deleted, but is stored in the recycle bin of the OBS file system instead. This section describes how to set a lifecycle rule for the recycle bin directory to periodically clear related data.

CAUTION

- For clusters that use decoupled storage and compute, configure a lifecycle policy for the related directories by referring to this chapter. Otherwise, the storage space may be used up and storage fees may increase. For details about OBS billing, see [OBS Billing Overview](#).
- The recycle bin directory is created per user. When a user is created in the MRS cluster and the user has the permission to delete component data, you need to configure the recycle bin clearing rule for this new user.
- For HBase components that use decoupled storage and compute in MRS 3.1.2 or later versions, refer to this topic to set a policy for clearing component data in the recycle bin.

You need to configure lifecycle policies for the recycle bin directories of preset users in the MRS cluster and the recycle bin directories of new users who need accidental deletion prevention. If a low privileged agency is used or only the permission for MRS users to access OBS file system directories is configured by referring to [Configuring Fine-Grained OBS Access Permissions for MRS Cluster Users](#), you will need the operation permission for the recycle bin directory.

Table 8-1 Directories for which a lifecycle policy needs to be configured

Cluster Version	Directory Type	Component	Directory	How to Create
Versions earlier than MRS 3.3.0-LTS	Recycle bin directories that must be configured by default for each component in an MRS cluster	Hive	<ul style="list-style-type: none"> user/omm/.Trash user/hive/.Trash 	If the .Trash folder does not exist, create it on the cluster client as user omm . Run the following command: hdfs dfs -mkdir -p obs://Name of the OBS parallel file system where the table is stored/ Folder path
		Spark	<ul style="list-style-type: none"> user/omm/.Trash user/root/.Trash user/spark2x/.Trash 	
		HetuEngine	<ul style="list-style-type: none"> user/omm/.Trash user/hetuserver/.Trash 	
		HBase	<ul style="list-style-type: none"> user/hbase/.Trash user/omm/.Trash 	
	Recycle bin directories of users who need accidental deletion prevention	Hive/Spark/HetuEngine	user/<New service user>/.Trash	
MRS 3.3.0-LTS or later	Default recycle bin directories configured for each component in an MRS cluster	Hive/Spark/HetuEngine	/user/.Trash	

For example, if a new user in the cluster has the following permissions, you need to create a recycle bin directory clearing rule for the user in the parallel file system:

- Permissions to delete the HDFS files
- **DROP, INSERT OVERWRITE, and TRUNCATE** permissions on Hive tables
- **DROP, TRUNCATE, DELETE, INSERT OVERWRITE, and LOAD OVERWRITE** permissions on HetuEngine

Configuring the Lifecycle Rule of an OBS Directory

- Step 1** Log in to the OBS console.
- Step 2** Click **Parallel File Systems** and click the name of the file system used by the current MRS cluster.
- Step 3** In the navigation pane on the left, choose **Basic Configurations > Lifecycle Rules**. Click **Create** to create a lifecycle rule for a specified directory. For details about the parameters, see [Configuring a Lifecycle Rule](#).

Table 8-2 Parameters for creating a lifecycle rule

Parameter	Description	Example Value
Status	Whether to enable the lifecycle rule.	Enable
Rule Name	Rule name that identifies different lifecycle configurations.	rule-test
Prefix	Prefix of the objects to which the lifecycle rule applies. Objects that have the specified prefix will be managed by the lifecycle rule. The prefix cannot start with a slash (/), have consecutive slashes (/), or contain the following special characters: \:*? "<> If this parameter is not specified, the rule will take effect for the entire file system. NOTE To prevent other service data from being deleted by mistake, you are not advised to use the lifecycle rule configured for the entire file system or high-level directories. Generally, the recycle bin directory of MRS components is in the following format. If the folder does not exist, create it. user/<Username>/.Trash	user/omm/.Trash
Delete Files After (Days)	The object within the rule configuration scope expires and is automatically deleted by OBS if the number of days since its last update reaches this parameter value.	30 days

- Step 4** Click **OK** to complete the lifecycle rule configuration.

You can click **Edit** in the **Operation** column of a lifecycle rule to edit it. You can also click **Disable** or **Enable** to disable or enable it.

- Step 5** Repeat the preceding steps to create recycle bin directory clearing rules for all users who have the data deletion permission in the current MRS cluster one by one until all recycle bin directories in the OBS file system are configured.

----End

8.2.3 Example for Interconnecting a Cluster Service with OBS

8.2.3.1 Interconnecting Flink with OBS Using an IAM Agency

After configuring decoupled storage and compute for a cluster by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you can access the OBS parallel file system using the Flink client and run jobs.

Interconnecting Flink with OBS

Step 1 Log in to the Flink client installation node as the client installation user.

Step 2 Initialize environment variables.

```
source Client installation directory/bigdata_env
```

Step 3 Configure the Flink client. For details, see [Using Flink from Scratch](#).

Step 4 Start a session.

- Normal cluster (Kerberos authentication disabled)

```
yarn-session.sh -nm "session-name" -d
```

- Security cluster (Kerberos authentication enabled)

- If the paths of the **flink.keystore** and **flink.truststore** files are relative ones:

Run the following command in the directory at the same level as **ssl** to start the session. **ssl/** is a relative path.

```
cd Client installation directory/Flink/flink/conf
```

```
yarn-session.sh -t ssl/ -nm "session-name" -d
```

```
...
```

```
Cluster started: Yarn cluster with application id application_1624937999496_0017  
JobManager Web Interface: http://192.168.1.150:32261
```

- If the paths of the **flink.keystore** and **flink.truststore** files are absolute ones:

Run the following command to start a session:

```
cd Client installation directory/Flink/flink/conf
```

```
yarn-session.sh -nm "session-name" -d
```

Step 5 Run the following command only on a security cluster with Kerberos authentication enabled to authenticate users:

```
kinit Username
```

Step 6 Explicitly add the OBS file system to be accessed in the Flink command line.

```
echo -e 'test' >/tmp/test
```

```
hdfs dfs -mkdir -p obs://Parallel file system name/tmp/flinkjob
```

```
hdfs dfs -put /tmp/test/ obs://Parallel file system name/tmp/flinkjob/
```

```
flink run Client installation directory/Flink/flink/examples/batch/WordCount.jar
-input obs://Parallel file system name/tmp/flinkjob/test -output obs://Parallel
file system name/tmp/flinkjob/output
```

----End

NOTE

Before interconnecting Flink with OBS, ensure that YARN is connected to OBS as Flink jobs run on YARN.

8.2.3.2 Interconnecting Flume with OBS Using an IAM Agency

After configuring decoupled storage and compute for a cluster by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you can run OBS jobs using Flume.

This section applies to MRS 3.x or later.

Interconnecting Flume with OBS

Step 1 Create an OBS folder for storing data.

1. Log in to the OBS console.
2. In the navigation pane on the left, choose **Resources** > **Parallel File Systems**.
3. On the displayed page, click the name of the parallel file system you created to access its details page.
4. In the navigation pane on the left, choose **Files**. On the displayed page, click **Create Folder** to create the **testFlumeOutput** folder.

Step 2 Log in to the node where the Flume client is installed as user **root**.

Step 3 Create the **/opt/flumeInput** directory and create a customized **.txt** file in it.

Step 4 Add the following content to the *Client installation directory*/**FusionInsight-flume-*/properties.properties** file:

```
# source
server.sources = r1
# channels
server.channels = c1
# sink
server.sinks = obs_sink
# ----- define net source -----
server.sources.r1.type = seq
server.sources.r1.spooldir = /opt/flumeInput
# ---- define OBS sink ----
server.sinks.obs_sink.type = hdfs
server.sinks.obs_sink.hdfs.path = obs://esdk-c-test-pfs1/testFlumeOutput
server.sinks.obs_sink.hdfs.filePrefix = %[localhost]
server.sinks.obs_sink.hdfs.useLocalTimeStamp = true
# set file size to trigger roll
server.sinks.obs_sink.hdfs.rollSize = 0
server.sinks.obs_sink.hdfs.rollCount = 0
server.sinks.obs_sink.hdfs.rollInterval = 5
#server.sinks.obs_sink.hdfs.threadPoolSize = 30
server.sinks.obs_sink.hdfs.fileType = DataStream
server.sinks.obs_sink.hdfs.writeFormat = Text
server.sinks.obs_sink.hdfs.fileCloseByEndEvent = false

# define channel
server.channels.c1.type = memory
```



```
server.channels.c1.capacity = 1000
# transaction size
server.channels.c1.transactionCapacity = 1000
server.channels.c1.byteCapacity = 800000
server.channels.c1.byteCapacityBufferPercentage = 20
server.channels.c1.keep-alive = 60
server.sources.r1.channels = c1
server.sinks.obs_sink.channel = c1
```

NOTE

- Set `server.sources.r1.spooldir` to the directory of the `.txt` file created in [Step 3](#).
- Set `server.sinks.obs_sink.hdfs.path` to the OBS file system created in [Step 1](#).

Step 5 Copy the `hadoop-huaweicloud-*.jar` and `mrs-obs-provider-*.jar` files from *Client installation directory/Hive/Beeline/lib* to *Flume client installation directory/fusionInsight-flume-*/lib*. Then run the following commands to modify permissions:

```
cd Flume client installation directory/fusionInsight-flume-*/lib
```

```
chmod 755 hadoop-huaweicloud-*.jar
```

```
chmod 755 mrs-obs-provider-*.jar
```

Step 6 Run the following command to restart the Flume client:

```
cd Flume client installation directory/fusionInsight-flume-*/bin
```

```
./flume-manager.sh restart
```

Step 7 View the result in the OBS system.

1. Log in to the OBS console.
2. In the navigation pane on the left, choose **Resources > Parallel File Systems**. Click the name of the parallel file system you created. In the navigation pane on the left, choose **Files**. On the displayed page, click the folder created in [Step 1](#) to view the result.

----End

8.2.3.3 Interconnecting HDFS with OBS Using an IAM Agency

After configuring decoupled storage and compute for a cluster by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you can view and create OBS file directories on the HDFS client.

Interconnecting HDFS with OBS

Step 1 Log in to the node where the HDFS client is installed as the client installation user.

Step 2 Run the following command to switch to the client installation directory.

```
cd Client installation directory
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 If the cluster is in security mode, authenticate the user. In normal mode, skip user authentication.

kinit *Component service user*

Step 5 Explicitly add the OBS file system to be accessed in the HDFS command line.

For example:

- Run the following command to access the OBS file system:

```
hdfs dfs -ls obs://OBS_parallel_file_system_name/Path
```

For example, run the following command to access the **mrs-word001** parallel file system. If the file list is returned, OBS is successfully accessed.

```
hadoop fs -ls obs://mrs-word001/
```

Figure 8-5 Returned file list

```
Found 2 items
drwxrwxrwx - root root          0 2019-12-21 11:04 obs://mrs-word001/input
drwxrwxrwx - root root          0 2019-12-21 11:04 obs://mrs-word001/program
```

- Run the following command to upload the **/opt/test.txt** file from the client node to the OBS file system path:

```
hdfs dfs -put /opt/test.txt obs://OBS_parallel_file_system_name/Path
```

----End

NOTE

If a large number of logs are printed in the OBS file system, the read and write performance may be affected. You can adjust the log level of the OBS client as follows:

```
cd Client_installation_directory/HDFS/hadoop/etc/hadoop
```

```
vi log4j.properties
```

Add the OBS log level configuration to the file as follows:

```
log4j.logger.org.apache.hadoop.fs.obs=WARN
```

```
log4j.logger.com.obs=WARN
```

Run the following command to view the log level:

```
tail -4 log4j.properties
```

Figure 8-6 Viewing the log level

```
[root@ecs-... hadoop]# tail -4 log4j.properties
# Log levels of third-party libraries
log4j.logger.org.apache.commons.beanutils=WARN
log4j.logger.org.apache.hadoop.fs.obs=WARN
log4j.logger.com.obs=WARN
[root@ecs-... hadoop]#
```

8.2.3.4 Interconnecting Hive with OBS Using an IAM Agency

After configuring decoupled storage and compute for a cluster by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you can create tables with OBS paths as their location on the Hive client.

Setting the Location to an OBS Path When Creating a Table

Step 1 Log in to the client installation node as the client installation user.

Step 2 Run the following command to initialize environment variables:

```
source Client installation directory/bigdata_env
```

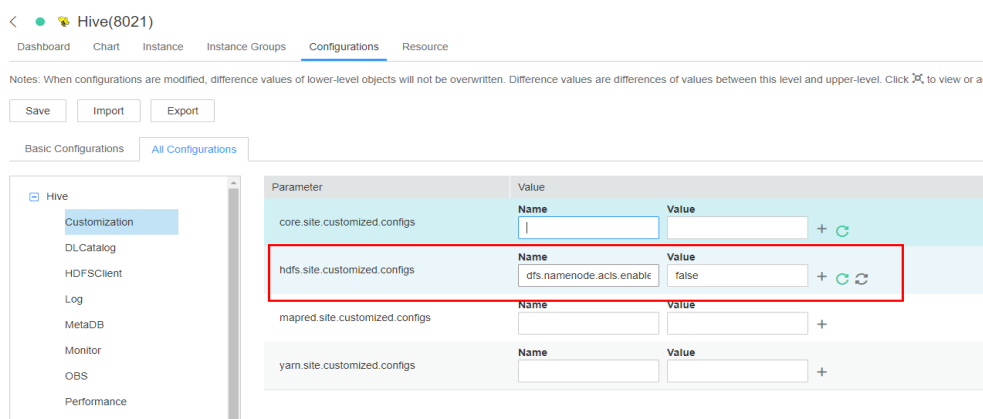
Step 3 For a security cluster, run the following command to perform user authentication (the user must have the permission to perform Hive operations). If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

```
kinit User performing Hive operations
```

Step 4 Log in to FusionInsight Manager of a cluster earlier than MRS 3.2.0, choose **Cluster > Services > Hive**, and click **Configurations > All Configurations**.

In the navigation pane on the left, choose **Hive > Customization**. In custom configuration items, add **dfs.namenode.acls.enabled** to **hdfs.site.customized.configs** and set its value to **false**.

Figure 8-7 Adding custom parameters



Step 5 Click **Save** to save the configuration for versions earlier than MRS 3.2.0. On the **Dashboard** page, click **More** and select **Restart Service**. Enter the password of the current user, click **OK**, and select **Restart upper-layer services**. Click **OK** to restart Hive.

Step 6 Log in to the beeline client and set **Location** to the OBS file system path when creating a table.

beeline

For example, run the following command to create the table **test** in **obs://OBS parallel file system name/user/hive/warehouse/Database name/Table name**.

```
create table test(name string) location "obs://OBS parallel file system name/  
user/hive/warehouse/Database name/Table name";
```

NOTE

You need to add the component operator to the URL policy in the Ranger policy. Set the URL to the complete path of the object on OBS. Select the Read and Write permissions.

For versions earlier than MRS 3.x, see [Configuring Hive Access Permissions in Ranger](#). For MRS 3.x or later, see [Adding a Ranger Access Permission Policy for Hive](#).


----End

Interconnecting Hive with OBS Through MetaStore

Step 1 Log in to FusionInsight Manager and choose **Cluster > Services > Hive > Configurations > All Configurations**.

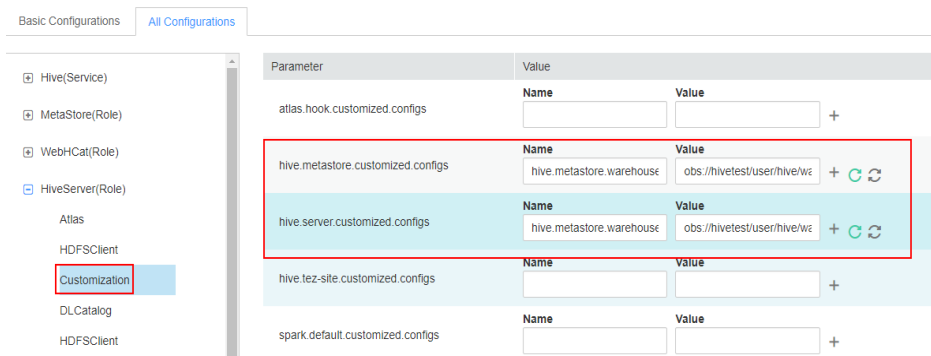
- For versions earlier than MRS 3.2.0:
 - In the navigation pane on the left, choose **MetaStore (role) > Customization**. Add the configuration item **hive.metastore.warehouse.dir** to the custom parameter **hive.metastore.customized.configs** and set the value to an OBS path. For example, set it to **obs://hivetest/user/hive/warehouse/**, where **hivetest** is the name of the OBS parallel file system.

Figure 8-8 Configuring **hive.metastore.warehouse.dir**




- In the navigation pane on the left, choose **HiveServer (role) > Customization**. Add the configuration item **hive.metastore.warehouse.dir** to **hive.metastore.customized.configs** and **hive.server.customized.configs** and set the value to an OBS path. For example, set it to **obs://hivetest/user/hive/warehouse/**, where **hivetest** is the name of the OBS parallel file system.

Figure 8-9 Configuring **hive.metastore.warehouse.dir**



- For MRS 3.2.0 and later versions:
Search for **hive.metastore.warehouse.dir** in the search box and change the parameter value to an OBS path, for example, **obs://hivetest/user/hive/warehouse/**. **hivetest** indicates the OBS file system name.

Figure 8-10 Configuring **hive.metastore.warehouse.dir**



- Step 2** Save the change and restart Hive.
- Step 3** (Optional) Install the client by referring to [Installing an MRS Cluster Client](#). If the client has been installed in the cluster, go to [Step 4](#).
- Step 4** Update the client configuration file.
- Run the following command to open `hivemetastore-site.xml` in the Hive configuration file directory on the client:
`vim Client installation directory/Hive/config/hivemetastore-site.xml`
 - Change the value of `hive.metastore.warehouse.dir` to the corresponding OBS path, for example, `obs://hivetest/user/hive/warehouse/`, where `hivetest` is the OBS bucket name.

Figure 8-11 Configuring the OBS Path

```

</property>
<property>
<name>hive.metastore.warehouse.dir</name>
<value>obs://hivetest/user/hive/warehouse</value>
</property>
</property>

```

- For MRS 3.2.0 and later versions, change the value of `hive.metastore.warehouse.dir` in `hivemetastore-site.xml` to the corresponding OBS path, for example, `obs://hivetest/user/hive/warehouse/`. The XML file is stored in the HCatalog client configuration file directory.
`vi Client installation directory/Hive/HCatalog/conf/hivemetastore-site.xml`
- Step 5** Log in to the beeline client, create a table, and check whether the location is the OBS path.

beeline

create table test(name string);

desc formatted test;

Location of the table is the OBS path.

Figure 8-12 Location of the Hive table

```

+-----+-----+
|                               data_type                               |
+-----+-----+
| data_type                     |
| string                        |
| NULL                          |
| NULL                          |
| default                       |
| USER                          |
| root                           |
| Wed May 10 19:18:31 CST 2023  |
| UNKNOWN                       |
| 0                               |
| obs://                        |
| MANAGED_TABLE                 |
| NULL                          |
| bucketing_version             |
| transient_lastDdlTime         |

```

NOTE

If the location of the current database points to HDFS, tables created in the database also point to HDFS by default. You do not need to specify the location. To modify the default table creation policy, modify the location of the database to point to OBS. Perform the following steps to modify the parameters:

1. Run the following command to query the location of the database:

```
show create database obs_test;
```

Figure 8-13 Viewing the location of the Hive Table

```
INFO : Concurrency mode is disabled, not creating a lock manager
+-----+
|          createdb_stmt          |
+-----+
| CREATE DATABASE `obs_test`      |
| LOCATION                        |
| 'hdfs://hacluster/user/hive/warehouse/obs_test.db' |
+-----+
3 rows selected (0.038 seconds)
```

2. Run the following command to change the database location:

```
alter database obs_test set location 'obs://OBS parallel file system name/user/hive/warehouse/Database name'
```

Run the **show create database** *obs_test* command to check whether the database location points to OBS.

Figure 8-14 Check the location of the modified Hive table.

```
INFO : Concurrency mode is disabled, not creating
+-----+
|          createdb_stmt          |
+-----+
| CREATE DATABASE `obs_test`      |
| LOCATION                        |
| 'obs://test1231/'              |
+-----+
3 rows selected (0.063 seconds)
```

3. Run the following command to modify the table location:

```
alter table user_info set location 'obs://OBS parallel file system name/user/hive/warehouse/Database name/Table name'
```

If the table contains data, migrate the original data file to the new location.

----End

8.2.3.5 Interconnecting Hudi with OBS Using an IAM Agency

After configuring decoupled storage and compute for a cluster by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you can create Hudi COW tables in spark-shell and store them to OBS.

Interconnecting Hudi with OBS

Step 1 Log in to the client installation node as the client installation user.

Step 2 Run the following commands to configure environment variables:

```
source Client installation directory/bigdata_env
```

```
source Client installation directory/Hudi/component_env
```

Step 3 Modify the configuration file:

```
vim Client installation directory/Hudi/hudi/conf/hdfs-site.xml
```

```
<property>  
<name>dfs.namenode.acls.enabled</name>  
<value>>false</value>  
</property>
```

Step 4 For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

```
kinit Username
```

Step 5 Start spark-shell and run the following commands to create a COW table and save it in OBS:

```
import org.apache.hudi.QuickstartUtils._  
import scala.collection.JavaConversions._  
import org.apache.spark.sql.SaveMode._  
import org.apache.hudi.DataSourceReadOptions._  
import org.apache.hudi.DataSourceWriteOptions._  
import org.apache.hudi.config.HoodieWriteConfig._  
val tableName = "hudi_cow_table"  
val basePath = "obs://testhudi/cow_table/"  
val dataGen = new DataGenerator  
val inserts = convertToStringList(dataGen.generateInserts(10))  
val df = spark.read.json(spark.sparkContext.parallelize(inserts, 2))  
df.write.format("org.apache.hudi").  
options(getQuickstartWriteConfigs).  
option(PRECOMBINE_FIELD_OPT_KEY, "ts").  
option(RECORDKEY_FIELD_OPT_KEY, "uuid").  
option(PARTITIONPATH_FIELD_OPT_KEY, "partitionpath").  
option(TABLE_NAME, tableName).  
mode(Overwrite).  
save(basePath);
```

 NOTE

"obs://*testhudi*/cow_table/" is the OBS path, and *testhudi* is the name of the parallel file system. Change them based on site requirements.

Step 6 Use DataSource to check whether the table is created and whether the data is normal.

```
val roViewDF = spark.
read.
format("org.apache.hudi").
load(basePath + "/*/*/*/*")
roViewDF.createOrReplaceTempView("hudi_ro_table")
spark.sql("select * from hudi_ro_table").show()
```

Figure 8-15 Viewing table data

_hoodie_commit_time	_hoodie_commit_seqno	_hoodie_record_key	_hoodie_partition_path	_hoodie_file_name	begin_lat	begin_lon	driver	end_lat	end_lon	fare	partitionpath	rider
202305162040091	20230516204009_1_314631e932-0c11-40e...	americas/united_s...l142858512594661971	americas/united_s...l0.2162415036760113616	americas/united_s...l9093714631e932-0c11-40e...	0.58909496240137841	0.09668238319271151	93.56818115236618	americas/united_s...l	rider-213116042065			
202305162040091	20230516204009_1_614946e241-5f5d-4b0...	americas/united_s...l0.62732122824896611	americas/united_s...l0.114883931570802611	americas/united_s...l2243214946e241-5f5d-4b0...	0.74546785375112951	0.39549398649089731	27.79478688582596	americas/united_s...l	rider-213116038282			
202305162040091	20230516204009_1_810c8f584a-d8d5-40c...	americas/united_s...l0.75202681532495821	americas/united_s...l0.87420415264005871	americas/united_s...l9591110e0f584a-d8d5-40c...	0.91970271200803821	0.362464778074404119	179139106643687	americas/united_s...l	rider-213116041640			
202305162040091	20230516204009_1_91a8609706-0af8-476...	americas/united_s...l0.49234796529128241	americas/united_s...l0.57318354079306341	americas/united_s...l238481a8609706-0af8-476...	0.8098858178093821610	0.42520899630713661	64.27696295804016	americas/united_s...l	rider-213116036812			
202305162040091	20230516204009_1_1012fb4e40e-10a5-4d7...	americas/united_s...l0.96945864170483921	americas/united_s...l0.18564880850682721	americas/united_s...l3649912fb4e40e-10a5-4d7...	0.3810636703720197410	0.252526522144790431	33.92216403948643	americas/united_s...l	rider-213116041375			
202305162040091	20230516204009_0_21eb27db0e-d5a3-486...	americas/brazil/s...l0.46157858458654831	americas/brazil/s...l0.616401eb27db0e-d5a3-486...	americas/brazil/s...l0.8304410444445281	0.7540034070000581	0.9671159942810241134	158204716302045	americas/brazil/s...l	rider-213116040336			
202305162040091	20230516204009_0_4118d28fa8-3290-439...	americas/brazil/s...l19204118d28fa8-3290-439...	americas/brazil/s...l0.87794022954277521	americas/brazil/s...l605581ac77c162-d08c-446...	0.34078785859296821	0.58307981422936551	43.4923811219014	americas/brazil/s...l	rider-213116040376			
202305162040091	20230516204009_2_119270e4c6-9d50-482...	asia/india/chemai...l0.56440921390409591	asia/india/chemai...l0.61811419270e4c6-9d50-482...	asia/india/chemai...l0.81928666877142241	0.79870638494151710	0.02690359227182834117	8511352550911551	asia/india/chemai...l	rider-213116041939			
202305162040091	20230516204009_2_5144a595df-e111-409...	asia/india/chemai...l0.81928666877142241	asia/india/chemai...l0.33290144a595df-e111-409...		0.2071489600291446210	0.862240318958269871	41.062989290463681	asia/india/chemai...l	rider-213116042167			

Step 7 Run :q to exit the spark-shell CLI.

----End

8.2.3.6 Interconnecting MapReduce with OBS Using an IAM Agency

After configuring decoupled storage and compute for a cluster by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you need to add custom configurations to MapReduce by performing the operations in this section.

Interconnecting MapReduce with OBS

Step 1 Log in to the MRS management console. In the navigation pane on the left, choose **Clusters > Active Clusters**. On the displayed page, click the name of the cluster you created to access its details page.

Step 2 Choose **Components > MapReduce**. The **All Configurations** page is displayed. In the navigation tree on the left, choose **MapReduce > Customization**. In the customized configuration items, add the configuration item

mapreduce.jobhistory.always-scan-user-dir to **core-site.xml** and set its value to **true**.

Figure 8-16 Adding a custom parameter

Parameter	Value	Description	Parameter File				
mapred.core-site.customized.configs	<table border="1"><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td>mapreduce.jobhistory.always-scan-us</td><td>true</td></tr></tbody></table>	Name	Value	mapreduce.jobhistory.always-scan-us	true	[Desc] Add a user customized configuration at MapR...	core-site.xml
Name	Value						
mapreduce.jobhistory.always-scan-us	true						

Step 3 Save the configurations and restart the MapReduce service.

----End

8.2.3.7 Interconnecting Presto with OBS Using an IAM Agency

After configuring decoupled storage and compute for a cluster by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you can use **presto_cli.sh** client to create tables and store them in OBS.

Interconnecting Presto with OBS

- For clusters with Kerberos authentication disabled
 - a. Log in to the node where the client is installed as the client installation user.
 - b. Run the following command to configure environment variables:
cd *Client installation directory*
source bigdata_env
 - c. Run the following command to connect to the client:
presto_cli.sh
 - d. Run the following command to create a schema and set **location** to an OBS path:
CREATE SCHEMA hive.demo WITH (location = 'obs://mrs-word001/presto-demo002');
 - e. Create a table in the schema. The table data is stored in the OBS file system. The following is an example:
CREATE TABLE hive.demo.demo_table WITH (format = 'ORC') AS SELECT * FROM tpch.sf1.customer;

Figure 8-17 Return result

```
[root@node-master2mdc0 ~]# presto_cli.sh
--server http://192.168.3.66:7520
presto> CREATE SCHEMA hive.demo WITH (location = 'obs://mrs-word001/presto-demo/');
CREATE SCHEMA
presto> CREATE TABLE hive.demo.demo_table WITH (format = 'ORC') AS SELECT * FROM tpch.sf1.customer;
CREATE TABLE: 150000 rows

Query 20191221_033019_00001_ukfbz, FINISHED, 2 nodes
Splits: 42 total, 42 done (100.00%)
0:09 [150K rows, 0B] [16K rows/s, 0B/s]
```

- f. Run **exit** to exit the client.
- For clusters with Kerberos authentication enabled
 - a. Log in to Manager and create a role with the Hive Admin Privilege permissions, for example, **prestorole**. For how to create a role, see [Managing MRS Cluster Roles](#).

- b. Create a user that belongs to the Presto and Hive groups and bind the role created in **a** to the user, for example, **presto001**. For how to create a user, see [Creating an MRS Cluster User](#).
- c. Authenticate the user.
kinit presto001
- d. Download the user credential.
 - For versions earlier than MRS 3.x, on MRS Manager, choose **System** > **Manage User**. Locate the row containing the new user, click **More**, and select **Download Authentication Credential**.
 - On FusionInsight Manager for MRS 3.x or later,, choose **System** > **Permission** > **User**. Locate the row containing the new user, click **More**, and select **Download Authentication Credential**.
- e. Decompress the downloaded user credential file and save the obtained **krb5.conf** and **user.keytab** files to the client directory, for example, *Client installation directory/Presto/*.
- f. Run the following command to obtain a user principal:
klist -kt Client installation directory/Presto/user.keytab
- g. Run the following command to connect to the Presto server of the cluster:
presto_cli.sh --krb5-config-path {krb5.conf file path} --krb5-principal {User principal} --krb5-keytab-path {user.keytab file path} --user {Presto username}
 - *krb5.conf file path*: Replace it with the file path set in **e**, for example, *Client installation directory/Presto/krb5.conf*.
 - *user.keytab file path*: Replace it with the file path set in **e**, for example, *Client installation directory/Presto/user.keytab (/opt/Bigdata/client/Presto/user.keytab)*.
 - *User principal*: Replace it with the result returned in **f**.
 - *Presto username*: Replace it with the username created in **b**, for example, **presto001**.
- h. Run the following command to create a schema and set **location** to an OBS path:
CREATE SCHEMA hive.demo01 WITH (location = 'obs://mrs-word001/presto-demo002/');
- i. Create a table in the schema. The table data is stored in the OBS file system. The following is an example:
CREATE TABLE hive.demo01.demo_table WITH (format = 'ORC') AS SELECT * FROM tpch.sf1.customer;

Figure 8-18 Return result

```
[root@node-master2@h2 ~]# presto_cli.sh --krb5-config-path /opt/client/Presto/krb5.conf --krb5-principal presto001@8985C7_17f6_488c_6763_09c42999AA1.COM --krb5-keytab-path /opt/client/Presto/user.keytab --user presto001
[presto@presto001 ~]$ klist -kt /opt/client/Presto/user.keytab
Krb5 Creds:
Principal: presto001@8985C7_17f6_488c_6763_09c42999AA1.COM
Expiration: 2024-09-23T10:00:00.000Z
Renewable: 2024-09-23T10:00:00.000Z
Initial: 2024-09-23T10:00:00.000Z
Last: 2024-09-23T10:00:00.000Z
[presto@presto001 ~]$ presto_cli.sh --krb5-config-path /opt/client/Presto/krb5.conf --krb5-principal presto001@8985C7_17f6_488c_6763_09c42999AA1.COM --krb5-keytab-path /opt/client/Presto/user.keytab --user presto001
presto> CREATE SCHEMA hive.demo01 WITH (location = 'obs://mrs-word001/presto-demo002/');
CREATE SCHEMA
presto> CREATE TABLE hive.demo01.demo_table WITH (format = 'ORC') AS SELECT * FROM tpch.sf1.customer;
CREATE TABLE: 10000 rows
Query: 2024-09-23T10:00:00.000Z: [cpu]: FINISHED, 2 nodes
sql14: 42 total, 42 done (100.0%)
0:11 [100k rows, 0] [13.7k rows/s, 60/s]
```

- j. Run **exit** to exit the client.

8.2.3.8 Interconnecting Spark with OBS Using an IAM Agency

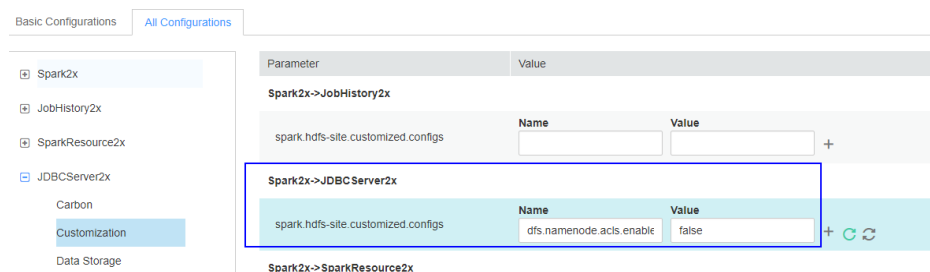
After configuring decoupled storage and compute for a cluster by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you can create tables with OBS paths as their location on the Spark client.

Verifying OBS Access with Spark Beeline

- Step 1** Log in to FusionInsight Manager and choose **Cluster > Services > Spark2x > Configurations > All Configurations**.


In the left navigation tree, choose **JDBCServer2x > Customization**. Add **dfs.namenode.acls.enabled** to the **spark.hdfs-site.customized.configs** parameter and set its value to **false**.

Figure 8-19 Adding Spark custom parameters



- Step 2** Search for the **spark.sql.statistics.fallBackToHdfs** parameter and set its value to **false**.

Figure 8-20 Setting **spark.sql.statistics.fallBackToHdfs**



- Step 3** Save the configurations and restart the JDBCServer2x instance.
- Step 4** Log in to the client installation node as the client installation user.
- Step 5** Run the following commands to configure environment variables:
- ```
source Client installation directory/bigdata_env
```
- Step 6** For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.
- ```
kinit Username
```
- Step 7** Access OBS using Spark beeline. The following example creates a table named **test** in the **obs://mrs-word001/table/** directory.
- ```
create table test(id int) location 'obs://mrs-word001/table/';
```

**Step 8** Run the following command to query all tables. If table **test** is returned, OBS access is successful.

**show tables;**

Figure 8-21 Returned table names

```
0: jdbc:hive2://ha-cluster/default> create table test(id int) location 'obs://mrs-word001/table/';
+-----+
| Result |
+-----+
No rows selected (2.515 seconds)
0: jdbc:hive2://ha-cluster/default> show tables;
+-----+
| database | tableName | isTemporary |
+-----+
| default | test | false |
| default | test_obs | false |
+-----+
2 rows selected (0.127 seconds)
```

**Step 9** Press **Ctrl+C** to exit Spark beeline.

----End

## Verifying OBS Access with Spark SQL

**Step 1** Log in to the client installation node as the client installation user.

**Step 2** Run the following commands to configure environment variables:

```
source Client installation directory/bigdata_env
```

**Step 3** Modify the configuration file:

```
vim Client installation directory/Spark2x/spark/conf/hdfs-site.xml
```

```
<property>
<name>dfs.namenode.acls.enabled</name>
<value>>false</value>
</property>
```

**Step 4** For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

```
kinit Username
```

**Step 5** Access OBS using Spark SQL CLI. For example, create a table named **test** in the **obs://mrs-word001/table/** directory.

1. Go to the **cd Client installation directory/Spark2x/spark/bin** directory and run the **./spark-sql** command to log in to the Spark SQL CLI.

2. Run the following command in the Spark SQL CLI:

```
create table test(id int) location 'obs://mrs-word001/table/';
```

**Step 6** Run the **show tables;** command to confirm that the table is created successfully.

**Step 7** Run **exit;** to exit the Spark SQL CLI.

**NOTE**

If a large number of logs are printed in the OBS file system, read and write performance may be affected. You can adjust the log level of the OBS client as follows:

```
cd Client installation directory/Spark2x/spark/conf
```


```
vi log4j.properties
```

Add the OBS log level configuration to the file as follows:

```
log4j.logger.org.apache.hadoop.fs.obs=WARN
```

```
log4j.logger.com.obs=WARN
```

Figure 8-22 Adding an OBS log level



```
[root@10-244-227-174 conf]#
[root@10-244-227-174 conf]# pwd
/opt/client_spark2x/Spark2x/spark/conf
[root@10-244-227-174 conf]# cat log4j.properties | grep obs
log4j.logger.org.apache.hadoop.fs.obs=WARN
log4j.logger.com.obs=WARN
[root@10-244-227-174 conf]#
```

----End

## Using Spark Shell to Read OBS Files

**Step 1** Log in to the client installation node as the client installation user.

**Step 2** Run the following commands to configure environment variables:

```
source Client installation directory/bigdata_env
```

**Step 3** Modify the configuration file:

```
vim Client installation directory/Spark2x/spark/conf/hdfs-site.xml
```

```
<property>
<name>dfs.namenode.acls.enabled</name>
<value>false</value>
</property>
```

**Step 4** For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

```
kinit Username
```

**Step 5** Create an OBS file.

1. Run the following commands to log in to the Spark SQL CLI:

```
cd Client installation directory/Spark2x/spark/bin
./spark-sql
```

2. Run the following commands to create a table and import data to the table:

```
create database test location "obs://Parallel file system path/test";
use test;
create table test1(a int,b int) using parquet;
insert into test1 values(1,2);
desc formatted test1;
```

Figure 8-23 Checking the location of the table

```
spark-sql> desc formatted test1;
a int NULL
b int NULL

Detailed Table Information
Database test1
Table test1
Owner root
Created Time Tue Nov 21 18:35:48 CST 2023
Last Access UNKNOWN
Created By Spark : -315088
Type MANAGED
Provider parquet
Location obs:/// /test1/test1
Serde Library org.apache.hadoop.hive.q1.io.parquet.serde.ParquetHiveSerDe
InputFormat org.apache.hadoop.hive.q1.io.parquet.MapredParquetInputFormat
OutputFormat org.apache.hadoop.hive.q1.io.parquet.MapredParquetOutputFormat
Time taken: 0.235 seconds, Fetched 16 row(s)
spark-sql>
```

**Step 6** Run the following command to go to the Spark **bin** directory:

```
cd Client installation directory/Spark2x/spark/bin
```

Run `./spark-sql` to log in to the Spark SQL CLI.

**Step 7** In the Spark Shell CLI, run the following command to query the table created in [Step 5.2](#):

```
spark.read.format("parquet").load ("obs://Parallel file system path/test1").show();
```

Figure 8-24 Viewing table data

```
scala> spark.read.format("parquet").load("obs:/// /test1/test1").show();
ERROR StatusLogger Log4j2 could not find a logging implementation. Please add log4j-core to the classpath. Using SimpleLogger to
log to the console...
2023-11-21 18:38:23,351 | WARN | main | The enable mv value "null" is invalid. Using the default value "false" | org.apache.car
bondata.core.util.CarbonProperties.validateEnableMV(CarbonProperties.java:512)
2023-11-21 18:38:23,366 | WARN | main | The value "LOCALLOCK" configured for key carbon.lock.type is invalid for current file s
ystem. Use the default value HDFSLOCK instead. | org.apache.carbondata.core.util.CarbonProperties.validateAndConfigureLockType(C
arbonProperties.java:441)
+-----+
| a | b |
+-----+
| 11 | 21 |
+-----+
```

**Step 8** Run the `:quit` command to exit the Spark Shell CLI.

```
----End
```

### 8.2.3.9 Interconnecting Sqoop with OBS Using an IAM Agency

After connecting the Sqoop client to an OBS file system by referring to [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#), you can import tables from a relational database to OBS or export tables from OBS to a relational database on the Sqoop client.

#### Prerequisites

You need to download the MySQL driver package of the required version from the MySQL official website <https://downloads.mysql.com/archives/c-j/>, decompress the package, and upload it to the `Client installation directory/Sqoop/sqoop/lib` directory on the node where the Sqoop client is installed.

## Exporting Sqoop Table Data to MySQL by Running sqoop export

**Step 1** Log in to the node where the client is installed.

**Step 2** Run the following command to initialize environment variables:

```
source /opt/client/bigdata_env
```

**Step 3** Run the following command to perform operations on the Sqoop client:

```
sqoop export --connect jdbc:mysql://10.100.xxx.xxx:3306/test --username root
--password xxx --table component13 --export-dir hdfs://hacluster/user/hive/
warehouse/component_test3 --fields-terminated-by ',' -m 1
```

For details about more parameters, see [Common Sqoop Commands and Parameters](#).

**Table 8-3** Parameter descriptions

Parameter	Description
--connect	URL used to connect to JDBC, in the format of <b>jdbc:mysql://IP address of the MySQL database.MySQL port/Database name</b> .
--username	Username used to log in to the MySQL database.
-password	Password used to log in to the MySQL database. To avoid potential security risks, disable history command recording before executing a command that contains authentication password information.
-table <table-name>	Name of the MySQL table used to store exported data.
-export-dir <dir>	HDFS path where the Sqoop table to be exported is located.
--fields-terminated-by	Delimiter of the data to be exported, which must be the same as that in the HDFS data table to be exported.
-m or -num-mappers <n>	<i>n</i> (4 by default) maps are started to import data concurrently. Make sure that this value is not greater than the maximum number of maps in a cluster.
-direct	A fast mode for importing data into a relational database using a tool like MySQL's mysqlimport, which is faster than using a JDBC connection.
-update-key <col-name>	Followed by a condition column name. You can use this parameter to update the existing data in a relational database.

Parameter	Description
-update-mode <mode>	How data is updated. The options are <b>updateonly</b> and <b>allowinsert</b> . It can only be used when the record to be imported does not exist in the relational data table. For example, if there is a record with <b>id=1</b> in the HDFS data to be imported and there is already a record with <b>id=2</b> in the table, the update will fail.
-input-null-string <null-string>	(Optional) If unset, <b>null</b> will be used.
-input-null-non-string <null-string>	(Optional) If unset, <b>null</b> will be used.
-staging-table <staging-table-name>	This parameter creates a table with the same data structure as the target table, stores all data in it, and then writes the results to the target table through a single transaction.  The parameter ensures transaction security during the process of importing data into a relational database table. Multiple transactions may occur during the import process, and if one transaction fails, it can affect other transactions, resulting in errors or duplicate records in the imported data. This parameter helps to avoid such situations.
-clear-staging-table	This parameter allows you to clear the data in the staging table before running the import process if it is not empty.

----End

## Importing MySQL Data to a Hive Table by Running sqoop import

**Step 1** Log in to the node where the client is installed.

**Step 2** Run the following command to initialize environment variables:

```
source /opt/client/bigdata_env
```

**Step 3** Run the following command to perform operations on the Sqoop client:

```
sqoop import --connect jdbc:mysql://10.100.xxx.xxx:3306/test --username root
--password xxx --table component --hive-import --hive-table component_test2
--delete-target-dir --fields-terminated-by "," -m 1 --as-textfile
```

**Table 8-4** Parameter descriptions

Parameter	Description
--hive-import	Imports data from a relational database to MRS Hive.



Parameter	Description
--delete-target-dir	Deletes the existing target file (if any) from Hive and reimports the file.
-append	Appends data to an existing dataset in HDFS. Once used, Sqoop imports data to file in a temporary directory, renames the file, and moves the file to a formal directory. This helps avoid duplicate file names in the directory.
-as-avrodatafile	Imports data to an Avro file.
-as-sequencefile	Imports data to a sequence file.
-as-textfile	Imports data to a text file. Once the file is created, you can query data in Hive by running SQL statements.
-boundary-query <statement>	SQL statement used to query boundaries. Before importing data, run this SQL statement to obtain a result set. Then import the data in the result set. The statement is like <b>-boundary-query 'select id,creationdate from person where id = 3'</b> (importing the record id=3) or <b>select min(&lt;split-by&gt;), max(&lt;split-by&gt;) from &lt;table name&gt;</b> .  If a field's value in the SQL statement is a string, the error message "java.sql.SQLException: Invalid value for getLong()" is displayed.
- columns<col,col,col...>	Fields to be imported, in the format of <i>-columns id,username</i> .
-direct	A fast mode for importing data into a relational database using a tool like MySQL's mysqlimport, which is faster than using a JDBC connection.
-direct-split-size	Splits imported data streams into byte-sized chunks, especially when importing data from PostgreSQL using a direct connection. It can split a file that reaches a set size into several independent files.
-inline-lob-limit	Maximum value of an inline LOB.
-m or -num-mappers	<i>n</i> (4 by default) maps are started to import data concurrently. Make sure that this value is not greater than the maximum number of maps in a cluster.
-query, -e<statement>	Imports data from query results and requires the specification of <b>-target-dir</b> and <b>-hive-table</b> . When using this parameter, the query statement must include a <b>WHERE</b> clause that contains <b>\$CONDITIONS</b> .  Example: <b>-query'select * from person where \$CONDITIONS' -target-dir /user/hive/warehouse/person -hive-table person</b>

Parameter	Description
-split-by<column-name>	Column name of a table, which is used to split work units and is generally followed by a primary key ID.
-table <table-name>	Name of the relational database table from which data is obtained.
-target-dir <dir>	HDFS path.
-warehouse-dir <dir>	Directory used to store data to be imported, which must not be used together with <b>-target-dir</b> . It is applicable when data is imported to HDFS but not Hive directories.
-where	<b>WHERE</b> clause when data is imported from a relational database, for example, <b>-where 'id = 2'</b> .
-z,-compress	Compresses sequence, text, and Avro data files using the GZIP compression algorithm. Data is not compressed by default.
-compression-codec	Hadoop compression codec, with GZIP used by default.
-null-string <null-string>	String to be interpreted as null for string columns. If not set, <b>NULL</b> will be used.
-null-non-string<null-string>	String to be interpreted as null for non-string columns. If not set, <b>NULL</b> will be used.
-check-column (col)	Column for checking incremental data import, for example, <b>id</b> .
-incremental (mode) append or lastmodified	Incrementally imports data. <b>append</b> : appends records, for example, appending records that are greater than the value specified by <b>last-value</b> . <b>lastmodified</b> : appends data that is modified after the date specified by <b>last-value</b> .
-last-value (value)	Maximum value (greater than the specified value) of the column after the last import, which can be set as needed.

----End

## Example Sqoop Usage

- Importing MySQL data to HDFS by running sqoop import  


```
sqoop import --connect jdbc:mysql://10.100.231.134:3306/test --username root --password xxx --query 'SELECT * FROM component where $CONDITIONS and component_id ="MRS 1.0_002"' --target-dir /tmp/component_test --delete-target-dir --fields-terminated-by "," -m 1 --as-textfile
```

- Exporting OBS data to MySQL by running sqoop export  
**sqoop export --connect jdbc:mysql://10.100.231.134:3306/test --username root --password xxx --table component14 --export-dir obs://obs-file-bucket/xx/part-m-00000 --fields-terminated-by ',' -m 1**
- Importing MySQL data to OBS by running sqoop import  
**sqoop import --connect jdbc:mysql://10.100.231.134:3306/test --username root --password xxx --table component --target-dir obs://obs-file-bucket/xx --delete-target-dir --fields-terminated-by ',' -m 1 --as-textfile**
- Importing MySQL data to a Hive foreign table stored on OBS by running sqoop import  
**sqoop import --connect jdbc:mysql://10.100.231.134:3306/test --username root --password xxx --table component --hive-import --hive-table component\_test01 --fields-terminated-by ',' -m 1 --as-textfile**

## Missing MySQL Driver Package When Importing or Exporting Data

If the error message "Could not load db driver class: com.mysql.jdbc.Driver" is displayed when the **sqoop import** or **sqoop export** command is executed, as shown in [Figure 8-25](#), the MySQL driver package is missing. To address this issue, download the MySQL driver package from the MySQL official website (<https://downloads.mysql.com/archives/c-j/>), decompress it, upload it to the *Client installation directory/Sqoop/sqoop/lib* directory, and run **sqoop import** or **sqoop export** to import or export data.

Figure 8-25 Missing MySQL driver package error



```
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type org.slf4j.impl.Log4jLoggerFactory
09:32:28.283 [main] INFO org.apache.sqoop.Sqoop - Running Sqoop version: 1.4.7
09:32:28.234 [main] WARN org.apache.sqoop.tool.BaseSqoopTool - Setting your password on the command-line is insecure. Consider using -P instead.
09:32:28.321 [main] INFO org.apache.sqoop.manager.HiveManager - Preparing to use a MySQL streaming resultset.
09:32:28.325 [main] ERROR org.apache.sqoop.Sqoop - Got exception running Sqoop: java.lang.RuntimeException: Could not load db driver class: com.mysql.jdbc.Driver
java.lang.RuntimeException: Could not load db driver class: com.mysql.jdbc.Driver
 at org.apache.sqoop.manager.SqlManager.makeConnection(SqlManager.java:877)
 at org.apache.sqoop.manager.GenericJdbcManager.getConnection(GenericJdbcManager.java:61)
 at org.apache.sqoop.manager.CatalogQueryManager.listDatabases(CatalogQueryManager.java:59)
 at org.apache.sqoop.tool.ListDatabasesTool.run(ListDatabasesTool.java:51)
 at org.apache.sqoop.Sqoop.run(Sqoop.java:149)
 at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:76)
 at org.apache.sqoop.Sqoop.runSqoop(Sqoop.java:185)
 at org.apache.sqoop.Sqoop.runTool(Sqoop.java:236)
 at org.apache.sqoop.Sqoop.runTool(Sqoop.java:245)
 at org.apache.sqoop.Sqoop.main(Sqoop.java:254)
[root@test-node-master-1012:~]#
```

## 8.2.4 Configuring Fine-Grained OBS Access Permissions for MRS Cluster Users

When fine-grained permission control is enabled, you can configure OBS access permissions to implement access control on directories in OBS file systems.

### NOTE

This section does not apply to MRS 1.9.2.

This function enables you to control MRS users' access to OBS resources. For example, if you allow user group A to only access log files in a specified OBS file system, perform the following operations:

1. Configure an agency with OBS access permissions for an MRS cluster so that OBS can be accessed using the temporary AK/SK automatically obtained by the ECS.

2. Create a policy on the IAM console to allow access to log files in a specified OBS file system, and create an agency bound to the policy permission.
3. In the MRS cluster, bind the new agency to user group A so that user group A only has the permission to access log files in the specified OBS file system.

In the following scenarios, the username used for submitting jobs is an internal username so that MRS multi-user access to OBS is not supported.

- In a cluster with Kerberos authentication enabled, the built-in username of **spark-beeline** for submitting jobs is **spark**. In a cluster with Kerberos authentication disabled, the built-in username is **omm**.
- In a cluster with Kerberos authentication enabled, the built-in username of **hbase shell** for submitting jobs is **hbase**. In a cluster with Kerberos authentication disabled, the built-in username is **omm**.
- In a cluster with Kerberos authentication enabled, the built-in usernames of Presto for submitting jobs is **omm** and **hive**. In a cluster with Kerberos authentication disabled, the built-in username is **omm**. On the cluster details page of the MRS console, choose **Components > Presto > Service Configuration**, set **Type** to **All**, search for **hive.hdfs.impersonation.enabled**, and change its value to **true**. In this way, multiple MRS users can access OBS with fine-grained permissions.

## Prerequisites

- Fine-grained permission control has been enabled. For details about permissions management, see [Creating an IAM User and Granting MRS Permissions](#).
- You have a basic knowledge of [Cloud Service Delegation](#) and OBS fine-grained policies.

## Configuring an Agency with OBS Access Permissions for a Cluster

Follow instructions in [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#) to configure an agency with OBS access permissions.

The agency takes effect for all users (including internal users) and user groups in the cluster. To control the permissions of users and user groups in the cluster to access OBS, perform the following operations.

### NOTE

When you configure permissions on an OBS path, if the write permission is configured, you need to configure the corresponding recycle bin path. The default recycle bin path is `/user/${current.user}/.Trash/`, where `${current.user}` indicates the current username.

## Creating a Policy and an Agency on IAM

Create policies with different access permissions and bind the policies to the agency. For details, see [Reference: Creating a Policy and an Agency on IAM](#).

## Configuring OBS Permission Control Mapping

- Step 1** On the MRS management console, choose **Active Clusters** and click the cluster name.

**Step 2** In the **Basic Information** area on the **Dashboard** tab page, click **Manage** next to **OBS Permission Control**.

**Step 3** Click **Add Mapping** and set parameters according to [Table 8-5](#).


**Table 8-5** Adding an OBS permission control mapping

Parameter	Description
IAM Agency	Select the agency created in <a href="#">Step 2</a> .
Type	<ul style="list-style-type: none"> <li>• <b>User:</b> User-level mapping</li> <li>• <b>Group:</b> User group-level mapping</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• User-level mapping takes priority over user group-level mapping. If you select <b>Group</b>, you are advised to enter the primary group name in <b>MRS User (User Group)</b>.</li> <li>• Do not use the same username (group) in multiple mapping records.</li> </ul>

Parameter	Description
MRS User (User Group)	<p>Use commas (,) to separate multiple names of users or user groups.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• If OBS permission control is not configured for a user and no AK and SK are configured, the <b>OBS OperateAccess</b> permission in <b>MRS_ECS_DEFAULT_AGENCY</b> will be used for accessing OBS. You are advised not to bind the internal user of a component to an agency.</li><li>• If you need to configure an agency for the internal user of a component when submitting a job in the following scenarios, the requirements are as follows:<ul style="list-style-type: none"><li>– To control permissions on spark-beeline operations, configure the username <b>spark</b> for clusters with Kerberos authentication enabled and the username <b>omm</b> for clusters with Kerberos authentication disabled.</li><li>– To control permissions on HBase shell operations, configure the username <b>hbase</b> for clusters with Kerberos authentication enabled and the username <b>omm</b> for clusters with Kerberos authentication disabled.</li><li>– To control permissions on Presto operations, configure usernames <b>omm</b>, <b>hive</b>, and the one used to log in to the client for clusters with Kerberos authentication enabled, and configure usernames <b>omm</b> and the one used to log in to the client for clusters with Kerberos authentication disabled.</li><li>– If you want to use Hive to create tables in beeline mode, set the username to the internal user <b>hive</b>.</li></ul></li></ul>

**Step 4** Click **OK**.

**Step 5** Select **I agree to authorize the trust relationships between MRS Users (Groups) and IAM agencies**, and click **OK**. The mapping between the MRS user and OBS permission is added.

If  appears next to **OBS Permission Control** on the **Dashboard** tab page or the mapping table has been updated for OBS permission control, the mapping takes effect. It takes about 1 minute to for the mapping to take effect.

In the **Operation** column of the mapping list, you can edit or delete the added mapping.

 NOTE

- If OBS permission control is not configured for a user and no AK and SK are configured, the permissions owned by the agency configured for the cluster in the **Object Storage Service (OBS)** project will be used to access OBS.
- Regardless of whether OBS permission control is configured, AK/SK permission is used for accessing OBS once it is configured.
- Security Administrator permission is required to modify, create, or delete a mapping.
- To apply the mapping changes in spark-beeline, hive beeline, and Presto, you need to restart Spark, exit beeline and enter again, and restart Presto, respectively.

----End

## Component Access to OBS When OBS Permission Control Is Enabled

**Step 1** Log in to any node in a cluster as user **root** using the password set during cluster creation.

**Step 2** Run the following commands to set the environment variables:

```
cd Client installation directory
```

```
source Client installation directory/bigdata_env
```

**Step 3** If the Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If the Kerberos authentication is disabled for the current cluster, skip this step:

```
kinit MRS cluster user
```

Example:

```
kinit admin
```

**Step 4** If Kerberos authentication is disabled for the current cluster, run the following commands to log in as a user who belongs to the **supergroup** group. Replace **XXXX** with the username. For how to create a user, refer to [Creating an MRS Cluster User](#).

```
mkdir /home/XXXX
```

```
chown XXXX /home/XXXX
```

```
su - XXXX
```

**Step 5** To access OBS, you do not need to configure the AK, SK, and endpoint.

OBS path format: **obs://*OBS parallel file system name*/*XXX***

```
hadoop fs -ls "obs://obs-example/job/hadoop-mapreduce-examples-3.1.2.jarobs-example/job/hadoop-mapreduce-examples-3.1.2.jar"
```

 NOTE

- To delete files on OBS using `hadoop fs` commands, run **`hadoop fs -rm -skipTrash`**.
- If data import is not involved when a table is created using `spark-sql` and `spark-beeline`, OBS will not be accessed. That is, if you create a table in an OBS directory on which you do not have permission, the **CREATE TABLE** operation will still be successful, but the error message "**403 AccessDeniedException**" is displayed when you insert data.

----End

## Reference: Creating a Policy and an Agency on IAM

### Step 1 Create a policy on IAM.

1. Log in to the IAM console.
2. In the navigation pane on the left, choose **Permissions > Policies/Roles**. On the displayed page, click **Create Custom Policy**.
3. Set parameters according to [Table 8-6](#). Obtain the customized OBS policy samples that are frequently used by referring to [OBS Custom Policies](#).

**Table 8-6** Policy parameters

Parameter	Description
Policy Name	Only letters, digits, spaces, and special characters (-_.,) are allowed.
Scope	Select <b>Global services</b> , because OBS is a global service.
Policy View	Select <b>Visual editor</b> .



Parameter	Description
Policy Content	<ul style="list-style-type: none"> <li>- <b>Allow:</b> Select <b>Allow</b>.</li> <li>- <b>Select service:</b> Select <b>Object Storage Service (OBS)</b>.</li> <li>- <b>Select action:</b> Select <b>ReadWrite, ReadOnly, and ListOnly</b>.</li> <li>- <b>Resources under All:</b> Select <b>Specific</b> and set the following parameters: <ul style="list-style-type: none"> <li>▪ <b>object:</b> Select <b>Specify resource path</b> and click <b>Add Resource Path</b> to add a path. The <b>/tmp</b> directory is used as an example. For example, enter <i>obs_bucket_name/tmp/</i> and <i>obs_bucket_name/tmp/*</i>.</li> <li>▪ <b>bucket:</b> Select <b>Specify resource path</b>, click <b>Add Resource Path</b>, and enter <i>obs_bucket_name</i>.</li> </ul> </li> <li>- (Optional) Request condition: not added.</li> </ul>
Description	(Optional) Brief description about the policy.

 **NOTE**

If the data write operation of each component is implemented in **rename** mode, the permission to delete objects must be configured when data is written.

4. Click **OK** to save the policy.

**Step 2** Create an agency on IAM.

1. Log in to the IAM console.
2. Choose **Agencies**. On the displayed page, click **Create Agency**.
3. Set parameters according to [Table 8-7](#).

**Table 8-7** Agency parameters

Parameter	Description
Agency Name	Only letters, digits, spaces, and special characters (-_.,) are allowed.
Agency Type	Select <b>Common account</b> .

Parameter	Description
Delegated Account	Enter your cloud account, that is, the account you register using your mobile phone number. It cannot be a federated user or an IAM user created using your cloud account.
Validity Period	Set this parameter as required.
Description	(Optional) Brief description about the agency.
Permissions	<ol style="list-style-type: none"><li>1. In the <b>Project [Region]</b> column, locate the row where <b>OBS</b> is, click <b>Attach Policy</b>.</li><li>2. Select the policy created in <b>Step 1</b> to display it in <b>Selected Policies</b>.</li><li>3. Click <b>OK</b>.</li></ol>

4. Click **OK** to save the agency.

 **NOTE**

If you modify an agency and policies bound to it after using the agency to access OBS, the modification will take effect within 15 minutes.

----End

## 8.3 FAQ About Decoupled Storage and Compute

### 8.3.1 How Do I Read Encrypted OBS Data When Running an MRS Job?

In MRS 1.9.x encrypted data in OBS file systems can be used to run jobs, and the encrypted job running results can be stored in OBS file systems. Currently, data can be accessed only through an OBS protocol.

OBS supports data encryption and decryption using KMS keys. All encryption and decryption operations are performed on OBS, and keys are managed by DEW.

To use the OBS encryption function in MRS, you must have the KMS Administrator permissions and configure the following settings for the corresponding component:

**NOTE**

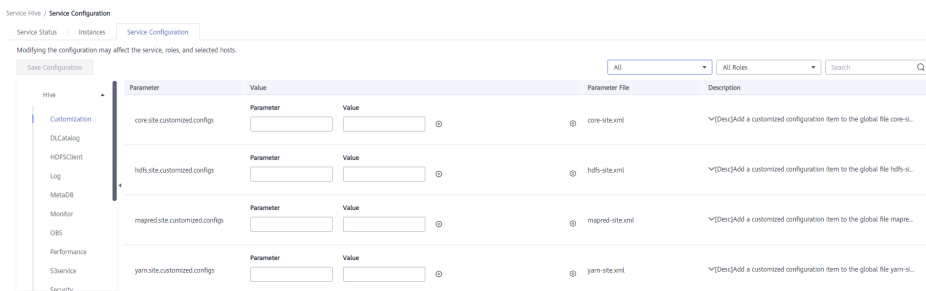
If the **OBS permission control** function is enabled in a cluster, the default agency **MRS\_ECS\_DEFAULT\_AGENCY** configured on the ECS or the AK/SK of the custom agency is used for accessing OBS. OBS uses the received AK/SK to access DEW to obtain the KMS key status. Therefore, you need to bind the KMS Administrator policy to the used agency. Otherwise, OBS returns the "403 Forbidden" error when processing encrypted data. Currently, the KMS Administrator policy is bound to the agency **MRS\_ECS\_DEFAULT\_AGENCY** by default. If you use a custom agency, you need to manually bind the policy to your custom agency.

**Prerequisites**

You have configured the function of accessing OBS from MRS first to use the OBS encryption function. For details, see [Interconnecting an MRS Cluster with OBS Using an IAM Agency](#).

**Hive Configuration**

- Step 1** Log in to the MRS console. On the **Active Clusters** page that is displayed, click the name of the desired cluster in the cluster list.
- Step 2** Choose **Components > Hive > Service Configuration**.
- Step 3** Switch **Basic** to **All**, and search for and set the following parameters:



**Table 8-8** Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	<ul style="list-style-type: none"> <li>● <b>SSE-KMS</b>: KMS keys are used for encryption and decryption</li> <li>● <b>NONE</b>: The encryption function is disabled.</li> </ul>
fs.obs.server-side-encryption-key	-	(Optional) This parameter indicates an ID of the KMS key used for encryption. If <b>fs.obs.server-side-encryption-type</b> is set to <b>SSE-KMS</b> and this parameter is not set, OBS uses the default KMS key for encryption.

Parameter	Value	Description
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. <ul style="list-style-type: none"> <li>• <b>true</b>: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to <b>true</b>.</li> <li>• <b>false</b>: The secure connection is disabled.</li> </ul>

**Step 4** Click **Save Configuration** and save the modified parameters as prompted.

----End

## Hadoop Configuration

### Method 1: Configuration on the GUI

**Step 1** Log in to the MRS console. On the **Active Clusters** page that is displayed, click the name of the desired cluster in the cluster list.

**Step 2** Choose **Components > HDFS > Service Configuration**.

**Step 3** Switch **Basic** to **All**, and search for and set the following parameters:

**Table 8-9** Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	<ul style="list-style-type: none"> <li>• <b>SSE-KMS</b>: KMS keys are used for encryption and decryption</li> <li>• <b>NONE</b>: The encryption function is disabled.</li> </ul>
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If <b>fs.obs.server-side-encryption-type</b> is set to <b>SSE-KMS</b> and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. <ul style="list-style-type: none"> <li>• <b>true</b>: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to <b>true</b>.</li> <li>• <b>false</b>: The secure connection is disabled.</li> </ul>

**Step 4** Click **Save Configuration** and operate as prompted.

**Step 5** Log in to the Master node using **root** password you set during cluster creation. If there are multiple Master nodes, log in to each one and repeat **Step 5** to **Step 7**.

**Step 6** Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

```
cd /opt/Bigdata/client
```

**Step 7** Run the following command to update client configurations and enter the username and password. The username is **admin** and the password is the password for user **admin** you set during cluster creation.

```
./ autoRefreshConfig.sh
```

----End

### Method 2: Configuration Through the Client Configuration File

Add the following parameter settings to the client configuration file, for example, **/opt/Bigdata/client/HDFS/hadoop/etc/hadoop/core-site.xml**, on the Master node. If the cluster has multiple Master nodes, log in to each Master node and perform this operation.

**Table 8-10** Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	<ul style="list-style-type: none"><li>• <b>SSE-KMS</b>: KMS keys are used for encryption and decryption</li><li>• <b>NONE</b>: The encryption function is disabled.</li></ul>
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional.  If <b>fs.obs.server-side-encryption-type</b> is set to <b>SSE-KMS</b> and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. <ul style="list-style-type: none"><li>• <b>true</b>: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to <b>true</b>.</li><li>• <b>false</b>: The secure connection is disabled.</li></ul>

## HBase Configuration

### Method 1: Configuration on the GUI

- Step 1** Log in to the MRS console. On the **Active Clusters** page that is displayed, click the name of the desired cluster in the cluster list.
- Step 2** Choose **Components > HBase > Service Configuration**.
- Step 3** Switch **Basic** to **All**, and search for and set the following parameters:

**Table 8-11** Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	<ul style="list-style-type: none"><li>• <b>SSE-KMS</b>: KMS keys are used for encryption and decryption</li><li>• <b>NONE</b>: The encryption function is disabled.</li></ul>
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If <b>fs.obs.server-side-encryption-type</b> is set to <b>SSE-KMS</b> and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. <ul style="list-style-type: none"><li>• <b>true</b>: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to <b>true</b>.</li><li>• <b>false</b>: The secure connection is disabled.</li></ul>

- Step 4** Click **Save Configuration** and operate as prompted.
- Step 5** Log in to the Master node as user **root**. The password is the password of user **root** you set when you create the cluster. If the cluster has multiple Master nodes, log in to each Master node and repeat **Step 5** to **Step 7**.
- Step 6** Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

```
cd /opt/Bigdata/client
```

- Step 7** Run the following command to update client configurations and enter the username and password. The username is **admin** and the password is the password for user **admin** you set during cluster creation.

```
./ autoRefreshConfig.sh
```

```
----End
```

#### Method 2: Configuration Through the Client Configuration File

Add the following parameter settings to the client configuration file, for example, **/opt/Bigdata/client/HBase/hbase/conf/core-site.xml**, on the Master

node. If the cluster has multiple Master nodes, log in to each Master node and perform this operation.

**Table 8-12** Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	<ul style="list-style-type: none"><li>• <b>SSE-KMS</b>: KMS keys are used for encryption and decryption</li><li>• <b>NONE</b>: The encryption function is disabled.</li></ul>
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If <b>fs.obs.server-side-encryption-type</b> is set to <b>SSE-KMS</b> and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. <ul style="list-style-type: none"><li>• <b>true</b>: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to <b>true</b>.</li><li>• <b>false</b>: The secure connection is disabled.</li></ul>

## Spark Configuration

### Method 1: Configuration on the GUI

**Step 1** Log in to the MRS console. On the **Active Clusters** page that is displayed, click the name of the desired cluster in the cluster list.

**Step 2** Choose **Components > Spark > Service Configuration**.

**Step 3** Switch **Basic** to **All**, and search for and set the following parameters:

**Table 8-13** Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	<ul style="list-style-type: none"><li>• <b>SSE-KMS</b>: KMS keys are used for encryption and decryption</li><li>• <b>NONE</b>: The encryption function is disabled.</li></ul>

Parameter	Value	Description
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If <b>fs.obs.server-side-encryption-type</b> is set to <b>SSE-KMS</b> and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. <ul style="list-style-type: none"> <li><b>true</b>: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to <b>true</b>.</li> <li><b>false</b>: The secure connection is disabled.</li> </ul>

**Step 4** Click **Save Configuration** and operate as prompted.

**Step 5** Log in to the Master node as user **root**. The password is the password of user **root** you set when you create the cluster. If the cluster has multiple Master nodes, log in to each Master node and repeat **Step 5** to **Step 7**.

**Step 6** Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

```
cd /opt/Bigdata/client
```

**Step 7** Run the following command to update client configurations and enter the username and password. The username is **admin** and the password is the password for user **admin** you set during cluster creation.

```
./autoRefreshConfig.sh
```

----End

### Method 2: Configuration Through the Client Configuration File

Add the following parameter settings to the client configuration file, for example, **/opt/Bigdata/client/Spark/spark/conf/core-site.xml**, on the Master node. If the cluster has multiple Master nodes, log in to each Master node and perform this operation.

**Table 8-14** Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	<ul style="list-style-type: none"> <li><b>SSE-KMS</b>: KMS keys are used for encryption and decryption</li> <li><b>NONE</b>: The encryption function is disabled.</li> </ul>



Parameter	Value	Description
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If <b>fs.obs.server-side-encryption-type</b> is set to <b>SSE-KMS</b> and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. <ul style="list-style-type: none"><li>• <b>true</b>: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to <b>true</b>.</li><li>• <b>false</b>: The secure connection is disabled.</li></ul>

## Presto Configuration

**Step 1** Log in to the MRS console. On the **Active Clusters** page that is displayed, click the name of the desired cluster in the cluster list.

**Step 2** Choose **Components > Presto > Service Configuration**.

**Step 3** Switch **Basic** to **All**, and search for and set the following parameters:

**Table 8-15** Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	<ul style="list-style-type: none"><li>• <b>SSE-KMS</b>: KMS keys are used for encryption and decryption</li><li>• <b>NONE</b>: The encryption function is disabled.</li></ul>
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If <b>fs.obs.server-side-encryption-type</b> is set to <b>SSE-KMS</b> and this parameter is not set, OBS uses the default KMS key for encryption.

Parameter	Value	Description
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. <ul style="list-style-type: none"><li>• <b>true</b>: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to <b>true</b>.</li><li>• <b>false</b>: The secure connection is disabled.</li></ul>

**Step 4** Click **Save Configuration** and operate as prompted.

----End

## 8.3.2 Example Application Development for Interconnecting HDFS with OBS

### Interconnection Principles

- The code for creating a `FileSystem` object in HDFS queries the corresponding implementation class based on the URI scheme. That is, the implementation classes provided by different underlying layers are configured in the HDFS configuration file. HDFS creates the corresponding implementation class based on **fs.AbstractFileSystem.%s.impl**. The following is an example:

```
Create a file system instance for the specified uri using the conf. The conf is used to find the class name that implements the file system. The conf is also passed to the file system for its configuration.
*
*@param uri URI of the file system
*@param conf Configuration for the file system
*
*@return Returns the file system for the given URI
*
*@throws UnsupportedOperationException file system for <code>uri</code> is not found
*/
public static AbstractFileSystem createFileSystem(URI uri, Configuration conf)
 throws UnsupportedOperationException {
 final String fsImplConf = String.format("fs.AbstractFileSystem.%s.impl", uri.getScheme());

 Class<?> clazz = conf.getClass(fsImplConf, null);
 if (clazz == null) {
 throw new UnsupportedOperationException(String.format(
 "%s=null: %s: %s",
 fsImplConf, NO_ABSTRACT_FS_ERROR, uri.getScheme()));
 }
 return (AbstractFileSystem) newInstance(clazz, uri, conf);
}
```

- In core-default of HDFS, corresponding implementation classes have been added for different URLs such as `adl`, `hdfs`, and `file`.

```
<property>
 <name>fs.AbstractFileSystem.adl.impl</name>
 <value>org.apache.hadoop.fs.adl.AdL</value>
</property>
<property>
 <name>fs.AbstractFileSystem.hdfs.impl</name>
 <value>org.apache.hadoop.fs.Hdfs</value>
 <description>The FileSystem for hdfs: uris.</description>
</property>
```

```
<property>
 <name>fs.AbstractFileSystem.file.impl</name>
 <value>org.apache.hadoop.fs.local.LocalFs</value>
 <description>The AbstractFileSystem for file: uris.</description>
</property>

<property>
 <name>fs.AbstractFileSystem.har.impl</name>
 <value>org.apache.hadoop.fs.HarFs</value>
 <description>The AbstractFileSystem for har: uris.</description>
</property>
```

- The OBS implementation class has been added to the default configuration file of MRS to connect to OBS.

```
<property>
 <name>fs.AbstractFileSystem.obs.impl</name>
 <value>org.apache.hadoop.fs.obs.OBS</value>
</property>
```

## Obtaining the Configuration File of a Cluster

1. Download and decompress the client by referring to [Installing an MRS Cluster Client](#).
2. Obtain **core-site.xml** and **hdfs-site.xml** from the downloaded HDFS client configuration file (*Download path/HDFS/hadoop/etc/hadoop*) and **core-site.xml** from the YARN client configuration file (*Download path/Yarn/config*).

### NOTE

These files are used to replace the configuration files used in the original code.

3. Add the following OBS access information to HDFS' and YARN's **core-site.xml** files:

```
<property>
 <name>fs.obs.endpoint</name>
 <value>obs endpoint</value>
</property>
<property>
 <name>fs.obs.access.key</name>
 <value>xxx</value>
 <description>huaweicloud access key</description>
</property>
<property>
 <name>fs.obs.secret.key</name>
 <value>xxx</value>
 <description>huaweicloud secret key</description>
</property>
```

### NOTE

- Configuration files containing authentication passwords pose security risks. Delete such files after configuration or store them securely.
4. Change the value of **fs.defaultFS** in the **core-site.xml** file on the HDFS client. For example, the value is **hdfs://hacluster** before the change.

```
<property>
 <name>fs.defaultFS</name>
 <value>hdfs://hacluster</value>
</property>
```

Change the value to **obs://Bucket name**.

```
<property>
 <name>fs.defaultFS</name>
 <value>obs://Bucket name</value>
</property>
```

- To reduce OBS logs, add the following configuration to the **log4j.properties** file:

```
log4j.logger.org.apache.hadoop.fs.obs=WARN
log4j.logger.com.obs=WARN
```

#### NOTE

If a large number of logs are printed in the OBS file system, the read and write performance may be affected. You can adjust the log level of the OBS client as follows:

```
cd ${client_home}/HDFS/hadoop/etc/hadoop
```

```
vi log4j.properties
```

Add the OBS log level configuration to the file as follows:

```
log4j.logger.org.apache.hadoop.fs.obs=WARN
log4j.logger.com.obs=WARN
```

```
[root@node-master1AuKK hadoop]# tail -4 log4j.properties
log4j.logger.org.apache.commons.beanutils=WARN

log4j.logger.org.apache.hadoop.fs.obs=WARN
log4j.logger.com.obs=WARN
[root@node-master1AuKK hadoop]#
```

## Adding Dependency Packages to Service Programs

Obtain the JAR files **hadoop-huaweicloud-xxx-hw-xx.jar** and **mrs-obs-provider-xxx.jar** from the MRS HDFS client installation package, place them in the **classpath** directory of the program, and modify the permissions and owner of the JAR files.

### 8.3.3 How Do I Connect an MRS Cluster Client to OBS Using an AK/SK Pair?

In MRS 1.9.2 or later, you can connect MRS clusters to OBS using **obs://**. Currently, supported components are Hadoop, Hive, Spark, Presto, and Flink. HBase cannot use **obs://** to interconnect with OBS.

This section describes how to use an AK/SK pair to connect MRS cluster components to OBS. The configuration file displays the AK/SK pair in plaintext. Be careful when utilizing the AK/SK pair.

**NOTICE**

- To improve data write performance, log in to the Manager and choose **Cluster > Services > Name of the service to be modified > Configurations**. Change the value of **fs.obs.buffer.dir** to the data disk directory.
- In big data storage-computing decoupling scenarios, make sure to use an OBS parallel file system. For details, see [Parallel File System](#). Using a regular object bucket can significantly impact the performance of the cluster.
- **In MRS 3.2.0-LTS.1 and later versions, components prevent mis-deletion by default. That is, file data deleted by component users is not directly deleted but stored in the recycle bin directory in the OBS file system.**  
To save OBS space, you need to enable periodical deletion of file data from the OBS recycle bin by referring to [Configuring the Policy for Clearing Recycle Bin Directories of MRS Cluster Components](#).
- Configuration files containing authentication passwords pose security risks. Delete such files after configuration or store them securely.
- Commands carrying authentication passwords pose security risks. Disable historical command recording before running such commands to prevent information leakage.

## Obtaining an AK/SK Pair and Endpoint

Obtain an AK, SK, and endpoint before interconnecting cluster components with OBS.

- Obtaining an AK/SK pair
  - a. Log in to the Huawei Cloud management console. Hover over the username in the upper right corner and select **My Credentials** from the drop-down list.
  - b. Click **Access Keys**. You can obtain the AK from the access key list and SK from the downloaded CSV file.

If there is no AK available, you can generate one by clicking **Add Access Key** and download it by entering the verification code or password and clicking **OK**.

**NOTE**

- You can only download the file before the page is closed. But, if you are unable to obtain the AK, you can recreate it.
  - To ensure security, keep your AK secure and update it regularly by deleting the old one and creating a new one.
- Obtaining an endpoint  
For how to obtain an endpoint, refer to [Regions and Endpoints](#).

## Using Hadoop to Access OBS

There are two ways to connect Hadoop to OBS. The first way is to add the AK/SK pair and endpoint to the **core-site.xml** file in the HDFS client installation directory. The second way is to add the AK/SK pair and endpoint when running Hadoop commands.

- Add the following content to the **core-site.xml** file in the *Client installation directory/HDFS/hadoop/etc/hadoop* directory on the HDFS client. For how to install the MRS client, see [Installing the Client](#).

```
<property>
 <name>fs.obs.access.key</name>
 <value>AK prepared in Obtaining an AK/SK Pair and Endpoint</value>
</property>
<property>
 <name>fs.obs.secret.key</name>
 <value>SK prepared in Obtaining an AK/SK Pair and Endpoint</value>
</property>
<property>
 <name>fs.obs.endpoint</name>
 <value>OBS endpoint prepared in Obtaining an AK/SK Pair and Endpoint</value>
</property>
```

If you use commands that need to submit jobs to YARN, such as **distcp**, add the preceding content to the **core-site.xml** file in the YARN directory (*Client installation directory/Yarn/config*) on the MRS client.

---

**NOTICE**

AK and SK will be displayed as plaintext in the configuration file. Exercise caution when setting AK and SK in the file.

---

Once the configuration is added, you can access data on OBS without the need to manually add the AK/SK and endpoint. For example, run the following command to view the file list of the **test\_obs\_orc** directory in the **obs-test** file system:

```
cd Client installation directory
```

```
source bigdata_env
```

```
kinit Component operation user (Skip this step if Kerberos authentication is disabled for the cluster.)
```

```
hadoop fs -ls "obs://obs-test/test_obs_orc"
```

- Add the AK/SK pair and endpoint to the command line to access data on OBS.

```
cd Client installation directory
```

```
source bigdata_env
```

```
kinit Component operation user (Skip this step if Kerberos authentication is disabled for the cluster.)
```

```
hadoop fs -Dfs.obs.endpoint=Endpoint prepared in Obtaining an AK/SK Pair and Endpoint -Dfs.obs.access.key=AK prepared in Obtaining an AK/SK Pair and Endpoint -Dfs.obs.secret.key=SK prepared in Obtaining an AK/SK Pair and Endpoint -ls "obs://obs-test/test_obs_orc"
```

## Using Hive to Access OBS

**Step 1** Log in to the service configuration page.


- For versions earlier than MRS 3.x, log in to the cluster details page and choose **Components > Hive > Service Configuration**.
- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing MRS Manager](#). Choose **Cluster > Services > Hive > Configurations**.

**Step 2** In the configuration type drop-down box, switch **Basic Configurations** to **All Configurations**.

**Step 3** Configure the AK and SK of OBS.


- For versions earlier than MRS 3.x, click **Hive**, select **Customization**, and add the following configurations to **core.site.customized.configs**: **Name: fs.obs.access.key**; **Value:** AK prepared in [Obtaining an AK/SK Pair and Endpoint](#); **Name: fs.obs.secret.key**; **Value:** SK prepared in [Obtaining an AK/SK Pair and Endpoint](#).

**Figure 8-26** Setting the AK/SK for accessing OBS



- For MRS 3.x or later, click **Hive(Service)**, select **Customization**, and add the following configurations to **core.site.customized.configs**: **Name: fs.obs.access.key**; **Value:** AK prepared in [Obtaining an AK/SK Pair and Endpoint](#); **Name: fs.obs.secret.key**; **Value:** SK prepared in [Obtaining an AK/SK Pair and Endpoint](#).

**Figure 8-27** Configuring the AK/SK for accessing OBS



**Step 4** Save the configurations and restart Hive.

**Step 5** Access the OBS directory in Beeline. For example, run the following command to create a Hive table and specify that data is stored in the **test\_obs** directory in the **test-bucket** file system:

```
cd Client installation directory
```

```
source bigdata_env
```

```
kinit Component operation user (Skip this step if Kerberos authentication is disabled for the cluster.)
```

```
create table test_obs(a int, b string) row format delimited fields terminated by "," stored as textfile location "obs://test-bucket/test_obs";
```

```
----End
```

## Using Spark to Access OBS

### NOTE

- SparkSQL depends on Hive. Therefore, when configuring OBS on Spark, you need to modify the OBS configuration used in [Using Hive to Access OBS](#).
- In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

- spark-beeline and spark-sql

You can use spark-beeline or spark-sql to log in to the Spark client and run the following commands to configure AK and SK information for accessing OBS:

```
set fs.obs.access.key=AK prepared in Obtaining an AK/SK Pair and Endpoint;
```

```
set fs.obs.secret.key=SK prepared in Obtaining an AK/SK Pair and Endpoint;
```

```
set fs.obs.endpoint=Endpoint prepared in Obtaining an AK/SK Pair and Endpoint;
```


- spark-beeline

The spark-beeline can access OBS by configuring service parameters on Manager. The procedure is as follows:

- Log in to the service configuration page.
  - For versions earlier than MRS 3.x, log in to the cluster details page and choose **Components > Spark > Service Configuration**.
  - For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing MRS Manager](#). Choose **Cluster > Services > Spark2x > Configurations**.
- In the configuration type drop-down box, switch **Basic Configurations** to **All Configurations**.
- Choose **JDBCServer > OBS** and set **fs.obs.access.key** and **fs.obs.secret.key**.

If the preceding two parameters do not exist in the current cluster, choose **JDBCServer > Customization** in the navigation pane on the left and add the AK and SK prepared in [Obtaining an AK/SK Pair and Endpoint](#) to **spark.core-site.customized.configs**.

Figure 8-28 Adding parameters for accessing OBS



- Save the configurations and restart Spark.
- Access OBS in **spark-beeline**. For example, access the **obs://obs-demo-input/table/** directory.  
**create table test(id int) location 'obs://obs-demo-input/table/';**



- spark-sql and spark-submit

Both spark-sql and spark-submit can access OBS if you add the following content to the **core-site.xml** configuration file in the *Client installation directory/Spark/spark/conf* directory:

```
<property>
 <name>fs.obs.access.key</name>
 <value>AK prepared in Obtaining an AK/SK Pair and Endpoint</value>
</property>
<property>
 <name>fs.obs.secret.key</name>
 <value>SK prepared in Obtaining an AK/SK Pair and Endpoint</value>
</property>
<property>
 <name>fs.obs.endpoint</name>
 <value>Endpoint prepared in Obtaining an AK/SK Pair and Endpoint</value>
</property>
```

## Using Presto to Access OBS

- Step 1** Go to the cluster details page and choose **Components > Presto > Service Configuration**.
- Step 2** In the configuration type drop-down box, switch **Basic Configurations** to **All Configurations**.
- Step 3** Search for and configure the following parameters:
- Set **fs.obs.access.key** to the AK prepared in [Obtaining an AK/SK Pair and Endpoint](#).
  - Set **fs.obs.secret.key** to the SK prepared in [Obtaining an AK/SK Pair and Endpoint](#).
- If the preceding two parameters cannot be found in the current cluster, choose **Presto > Hive** in the navigation tree on the left and add the two parameters to the customized parameter **core.site.customized.configs**.
- Step 4** Save the configurations and restart Presto.
- Step 5** Choose **Components > Hive > Service Configuration**.
- Step 6** In the configuration type drop-down box, switch **Basic Configurations** to **All Configurations**.
- Step 7** Search for and configure the following parameters:
- Set **fs.obs.access.key** to the AK prepared in [Obtaining an AK/SK Pair and Endpoint](#).
  - Set **fs.obs.secret.key** to the SK prepared in [Obtaining an AK/SK Pair and Endpoint](#).
- Step 8** Save the configurations and restart Hive.
- Step 9** On the Presto client, run the following statement to create a schema and set **location** to an OBS path:

```
presto_cli.sh
```

```
CREATE SCHEMA hive.demo WITH (location = 'obs://obs-demo/presto-demo/');
```

**Step 10** Create a table in the schema. The table data is stored in the OBS file system. The following is an example:

**USE** `hive.demo;`

**CREATE TABLE** `Table name (id int);`

**INSERT INTO** `Table name VALUES (2);` In this command, 2 is only used as an example. Replace it with the real value.

**CREATE TABLE** `hive.demo.demo_table WITH (format = 'ORC') AS SELECT * FROM tpch.sf1.customer;`

`----End`

## Using Flink to Access OBS

Add the following configuration to the Flink configuration file of the MRS client in *Client installation path/Flink/flink/conf/flink-conf.yaml*:

`fs.obs.access.key: AK prepared in Obtaining an AK/SK Pair and Endpoint`

`fs.obs.secret.key: SK prepared in Obtaining an AK/SK Pair and Endpoint`

`fs.obs.endpoint: Endpoint prepared in Obtaining an AK/SK Pair and Endpoint`

### NOTICE

AK and SK will be displayed as plaintext in the configuration file. Exercise caution when setting AK and SK in the file.

Once the configuration is added, you can access data on OBS without the need to manually add the AK/SK and endpoint.

## 8.3.4 How Do I Access OBS Using an MRS Client Installed Outside a Cluster?

### Scenario

In storage-compute decoupling scenarios where data is computed in an MRS cluster and stored in OBS buckets, you can obtain a temporary AK/SK using an agency and then use the AK/SK to access the OBS server. To access OBS from a client on a node outside the cluster, obtain an AK/SK through Guardian. Guardian is a component developed by the MRS team that allows clients outside the cluster to access OBS using temporary AKs/SKs.

### NOTE

This function is available in MRS 3.1.5 or later.

## How Does Guardian Enable a Client on a Node Outside the Cluster to Access OBS?

**Step 1** Install Guardian in the cluster.

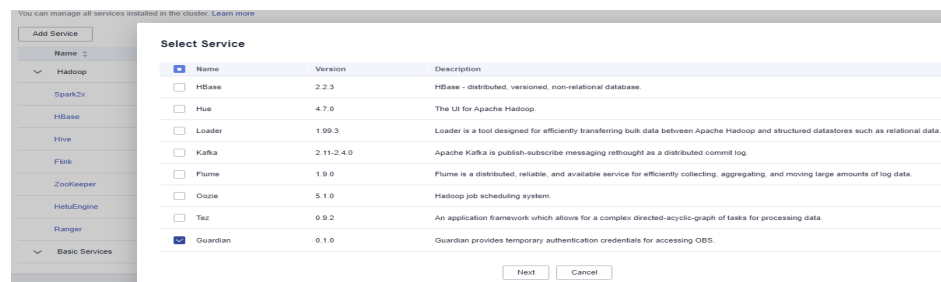
- For a new cluster in creation, select **Guardian**.

**Figure 8-29** Selecting Guardian

<input checked="" type="checkbox"/>	Zookeeper	3.6.3	A centralized service for maintaining configuration information, naming, performing distributed synchronization, and providing group services.
<input checked="" type="checkbox"/>	Ranger	2.0.0	RANGER is a framework to enable, monitor and manage comprehensive data security across the Hadoop platform.
<input type="checkbox"/>	Tez	0.9.2	An application framework which allows for a complex directed-acyclic-graph of tasks for processing data.
<input type="checkbox"/>	Impala	3.4.0	An SQL query engine for processing huge volumes of data.
<input type="checkbox"/>	Presto	333	An open source distributed SQL query engine.
<input type="checkbox"/>	ClickHouse	21.3.4.25	ClickHouse is a column-oriented database management system(DBMS) for online analytical processing of queries(OLAP).
<input type="checkbox"/>	Kudu	1.12.1	Kudu is a columnar storage manager developed for the Apache Hadoop platform.
<input type="checkbox"/>	Sqoop	1.4.7	Sqoop is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured datastores such as relational databases.
<input checked="" type="checkbox"/>	Guardian	0.1.0	Guardian provides temporary authentication credentials for accessing OBS.

- For an existing cluster, add the Guardian component in the **Components** tab.
  - a. On the cluster details page, choose **Components** and click **Add Service**.
  - b. In the service list, select the services to be added and click **Next**.

**Figure 8-30** Adding Guardian



- c. On the **Topology Adjustment** page, select the node where the service is to be deployed. (You are advised to deploy the Guardian service on the master node.)
- d. Click **OK**. After the service is added, you can view the added service on the **Components** page.

**NOTE**

The services added on the console are automatically synchronized to Manager.

**Step 2** Configure storage and compute decoupling for the cluster by referring to [Configuring a Storage-Compute Decoupled Cluster \(Agency\)](#) or [Configuring Fine-Grained Permissions for MRS Multi-User Access to OBS](#).

**Step 3** Install or update the client.

- For details about how to install the client on a node outside the cluster, see [Installing a Client \(MRS 3.x\)](#).
- For details about how to update an existing client, see [Updating the MRS Cluster Client After the Server Configuration Expires](#).

**Step 4** Once the installation is successful, you can access OBS. For example, access OBS from the HDFS client of an MRS cluster with Kerberos authentication disabled:

1. Log in to the node where the client is installed as the client installation user.
2. Go to the client installation directory and configure the environment variables:

```
cd Client installation directory
source bigdata_env
```

3. Run the following command on the HDFS client to access OBS:

```
hdfs dfs -ls obs://Directory of the OBS parallel file system
```

If information similar to the following is displayed, the interconnection is successful:

```
023-01-10 16:07:35 167|com.obs.services.AbstractClient|doActionWithResult|393|Storage|HTTP+XML|listObjects|||2023-01-10 16:07:35|2023-01-10 16:07:35|||0|
023-01-10 16:07:35 167|com.obs.services.AbstractClient|doActionWithResult|394|ObsClient |[listObjects] cost 185 ms

Found 2 items
rwkrwxrwx - root root 0 2022-09-07 15:10 obs:// /test
rwkrwxrwx - root root 0 2022-09-07 15:10 obs:// /user
023-01-10 16:07:35,174 INFO obs.OBSFileSystem: Finish closing filesystem instance for url: obs:// l
```

----End

## 8.3.5 Accessing an MRS Cluster's Manager (Version 2.x or Earlier)


### Scenario

MRS allows you to oversee, adjust, and handle clusters on Manager. Once the cluster is set up, you can access Manager as user **admin**.

Currently, you can access Manager using the following methods:

- **Accessing MRS Manager Using an EIP:** You can bind an EIP to the cluster to access the MRS Manager GUI and open source components managed in the cluster. This method is suggested as it is more user-friendly.
- **Accessing MRS Cluster Manager Using Direct Connect:** Direct Connect is a high-speed, low-latency, stable, and secure dedicated network connection that connects your local data center to an online cloud VPC. It extends online cloud services and existing IT facilities to build a flexible, scalable hybrid computing environment.

You need to ensure that Direct Connect is available, and the connection between the local data center and the online VPC has been established. For details, see [What Is Direct Connect?](#)

You can switch between EIP access and Direct Connect access on the MRS console as follows: Log in to the MRS console, click  next to **MRS Manager** on the **Dashboard** page of the target MRS cluster, and switch between the two access methods on the displayed page.

- **Accessing MRS Manager Through an ECS:** Access Manager through an ECS that is in the same VPC as the MRS cluster. This method is complex and is recommended when the EIP function is not supported.
- **Configuring an SSH Tunnel to Access MRS Manager:** Users and an MRS cluster are in different networks. As a result, an SSH tunnel needs to be created to send users' requests for accessing websites to the MRS cluster and dynamically forward them to the target websites.

### Prerequisites

Make sure the cluster is not in the starting, stopping, stopped, deleting, deleted, or frozen state before accessing MRS Manager.

### Accessing MRS Manager Using an EIP

**Step 1** Log in to the MRS management console.

**Step 2** In the navigation pane, choose **Active Clusters**. Click the target cluster name to access the cluster details page.

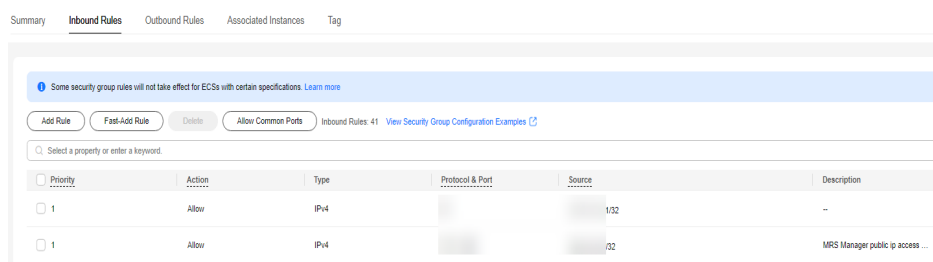
**Step 3** Click **Access Manager** next to **MRS Manager**. In the displayed dialog box, select **EIP** and configure the EIP information.

1. If no EIP is bound when the MRS cluster is created, select an available EIP from the EIP drop-down list. Otherwise, perform the operations in [Step 3.2](#).
  - If no available EIPs are displayed, click **Manage EIP** to create one. An EIP can be bound to only one MRS cluster.
  - To unbind or release an EIP, log in to the **EIPs** page, locate the row containing the target EIP, and click **Unbind** or choose **More > Release** in the **Operation** column.
2. In **Security Group**, select the security group to which the current cluster belongs. The security group is configured during cluster creation or is automatically created by the cluster.

If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

- "MRS Manager public ip access control rule" is automatically added to the **Description** column of the added security group. To view this description, choose **Manage Security Group Rule**, click **Security Group**, and click **Inbound Rules**.

**Figure 8-31** Adding a security group rule to the MRS cluster




- It is normal that the automatically generated public IP address is different from your local IP address and no action is required.
  - Port **9022** is the Knox port of the MRS cluster. Therefore, you need to enable the permission to access the port to access Manager.
3. Select the information to be confirmed and click **OK**. The Manager login page is displayed.

**Step 4** Enter the default username **admin** and the password set during cluster creation, and click **Log In**. The Manager page is displayed.

**Step 5** To grant users in other network segments the permission to access Manager, you can modify the security group and add the IP address range for the users to access the public network.

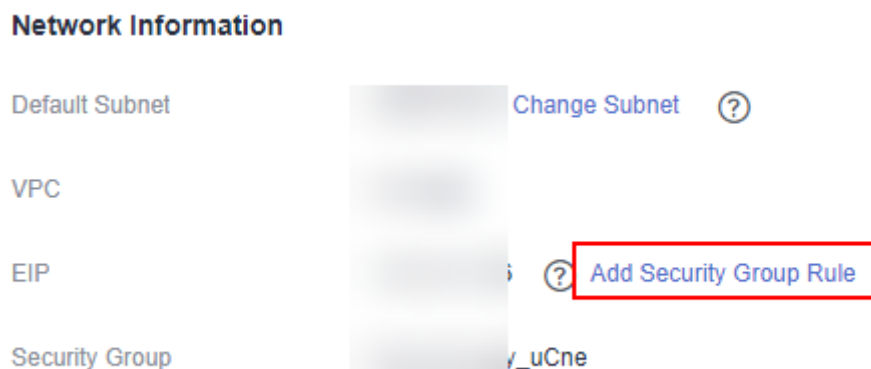
1. Click **Add Security Group Rule** next to **EIP**.

Figure 8-32 Cluster details



2. On the **Add Security Group Rule** page, add the IP address segment for users to access the public network and select **I confirm that *public network IP/port* is a trusted public IP address. I understand that using 0.0.0.0/0. poses security risks.** See [Figure 8-33](#).

Figure 8-33 Adding a security group rule



By default, the IP address segment used for accessing the public network is filled. You can change the IP address segment as required. To view, modify, or delete security group rules, click **Manage Security Group Rule**.

3. Click **OK**.

----End

## Accessing MRS Cluster Manager Using Direct Connect

- Step 1** Log in to the MRS console.
- Step 2** Click the name of the cluster to enter its details page.
- Step 3** On the **Dashboard** page of the cluster details page, click **Access Manager** next to **MRS Manager**.
- Step 4** Set **Access Mode** to **Direct Connect** and confirm that you understand the impact of the operation.

The floating IP address is automatically allocated by MRS to access MRS Manager. Before using Direct Connect to access MRS Manager, ensure that the connection between the local data center and the online VPC has been established.

**Step 5** Click **OK**. The MRS Manager login page is displayed. Enter the username **admin** and the password set during cluster creation.

----End

## Accessing MRS Manager Through an ECS

**Step 1** Log in to the MRS console.

**Step 2** On the **Active Clusters** page, click the name of the specified cluster.

Record the **AZ**, **VPC**, and **Security Group** of the cluster.

**Step 3** On the homepage of the management console, choose **Service List > Elastic Cloud Server** to switch to the ECS management console and create an ECS.

- The **AZ**, **VPC**, and **Security Group** of the ECS must be the same as those of the cluster to be accessed.
- Select a Windows public image. For example, a standard image **Windows Server 2012 R2 Standard 64bit(40GB)**.
- For details about other configuration parameters, see [Purchasing an ECS with Customized Configurations](#).

### NOTE

If the security group of the ECS is different from **Default Security Group** of the Master node, you can modify the configuration using either of the following methods:

- Change the security group of the ECS to the default security group of the Master node. For details, see [Changing a Security Group](#).
- Add two security group rules to the security groups of the Master and Core nodes to enable the ECS to access the cluster. Set **Protocol** to **TCP**, **Ports** of the two security group rules to **28443** and **20009**, respectively. For details, see [Creating a Security Group](#).

If "Failed to add security group rules." is displayed, check whether the security group quota is sufficient. If more quotas are needed, increase the quotas or delete security group rules that are no longer used.

**Step 4** On the EIP console, apply for an EIP and bind it to the ECS.

For details, see [Assigning an EIP and Binding It to an ECS](#).

**Step 5** Log in to the ECS.


The Windows system account, password, EIP, and security group rules are required for logging in to the ECS. For details, see [Logging In to a Windows ECS](#).

**Step 6** On the Windows remote desktop, use your browser to access Manager.

The Manager access address is in the **https://OMS floating IP address:28443/web** format. Enter the name and password of the cluster user, for example, user **admin**.

 NOTE

- To obtain the floating IP address of OMS, log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of OMS. Record the value of **inet**. If the floating IP address of OMS cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one Master node, query and record the cluster manager IP address of the Master node.
- If you access Manager with other cluster usernames, change the password upon your first access. The new password must meet the requirements of the current password complexity policies. For details, contact the administrator.
- By default, a user is locked after inputting an incorrect password five consecutive times. The user is automatically unlocked after 5 minutes.

**Step 7** Log out of FusionInsight Manager. To log out of Manager, move the cursor to  in the upper right corner and click **Log Out**.

----End

## Configuring an SSH Tunnel to Access MRS Manager

Users and an MRS cluster are in different networks. As a result, an SSH tunnel needs to be created to send users' requests for accessing websites to the MRS cluster and dynamically forward them to the target websites.

The MAC system does not support this function. For details about how to access MRS, see [Accessing MRS Manager Using an EIP](#).

Make sure the following prerequisites are met before proceeding with the operation:

- You have prepared an SSH client for creating the SSH tunnel, for example, the Git open source SSH client. You have downloaded and installed the client.
- You have created a cluster and prepared a key file in PEM format or obtained the password used during cluster creation.
- Users can access the Internet on the local PC.

**Step 1** Log in to the MRS console and click **Active Clusters**.

**Step 2** Click the specified MRS cluster name.

Record the security group of the cluster.

**Step 3** Add an inbound rule to the security group of the master node to allow data access to the IP address of the MRS cluster through port **22**.

For details, see [Adding a Security Group Rule](#).

**Step 4** Query the primary management node of the cluster by referring to [Checking MRS Active/Standby Management Nodes](#).

**Step 5** Bind an elastic IP address to the primary management node of the cluster.

For details, see [Assigning an EIP and Binding It to an ECS](#).

**Step 6** Start Git Bash locally and run the following command to log in to the active management node of the cluster:



```
ssh root@EIP address
```

Alternatively, run the following command:

```
ssh -i Key file path root@EIP address
```

**Step 7** View data forwarding configurations.

```
cat /etc/sysctl.conf | grep net.ipv4.ip_forward
```

- If **net.ipv4.ip\_forward=1** is displayed, the forwarding function has been configured. Go to [Step 9](#).
- If **net.ipv4.ip\_forward=0** is displayed, the forwarding function has not been configured. Go to [Step 8](#).
- If **net.ipv4.ip\_forward** fails to be queried, this parameter has not been configured. Run the following command and then go to [Step 9](#):  

```
echo "net.ipv4.ip_forward = 1" >> /etc/sysctl.conf
```

**Step 8** Modify forwarding configurations on the node.

1. Switch to user **root**.

```
sudo su - root
```

2. Modify forwarding configurations.

```
echo 1 > /proc/sys/net/ipv4/ip_forward
```

```
sed -i "s/net.ipv4.ip_forward=0/net.ipv4.ip_forward = 1/g" /etc/sysctl.conf
sysctl -w net.ipv4.ip_forward=1
```

3. Modify the **sshd** configuration file.

```
vi /etc/ssh/sshd_config
```

Press **I** to enter the edit mode. Locate **AllowTcpForwarding** and **GatewayPorts** and delete comment tags. Modify them as follows. Save the changes and exit.

```
AllowTcpForwarding yes
GatewayPorts yes
```

4. Restart the **sshd** service.

```
service sshd restart
```

**Step 9** View the floating IP address.

```
ifconfig
```

In the command output, **eth0:FI\_HUE** indicates the floating IP address of Hue, and **eth0:wsom** indicates the floating IP address of Manager. Record the value of **inet**.

Run the **exit** command to exit.

**Step 10** Run the following command on the local host to create an SSH tunnel that supports dynamic port forwarding:

```
ssh -i Path of the key file -v -ND Local port root@EIP address
```

Alternatively, run the following command:

```
ssh -v -ND Local port root@EIP address
```

Enter the password for creating the cluster as prompted.

In the command, set **Local port** to the user's local port that is not occupied. Port **8157** is recommended.

After the SSH tunnel is created, use **-D** to enable the dynamic port forwarding function. By default, the dynamic port forwarding function enables a SOCKS proxy process and monitors the user's local port. Port data will be forwarded to the primary management node using the SSH tunnel.

**Step 11** Configure the browser proxy.

1. Go to the Google Chrome client installation directory on the local PC.
2. Press **Shift** and right-click the blank area, choose **Open Command Window Here** and enter the following command:

```
chrome --proxy-server="socks5://localhost:8157" --host-resolver-rules="MAP * 0.0.0.0 , EXCLUDE localhost" --user-data-dir=c:/tmp/path --proxy-bypass-list="*google*.com,*gstatic.com,*gvt*.com,*:80"
```

 **NOTE**

- In the preceding command, **8157** is the local proxy port configured in [Step 10](#).
- If the local OS is Windows 10, start the Windows OS, click **Start** and enter **cmd**. In the displayed CLI, run the command in [Step 11.2](#). If this method fails, click **Start**, enter the command in the search box, and run the command in [Step 11.2](#).

**Step 12** In the address box of the browser, enter the address for accessing Manager.

The Manager access address is in the **https://Manager floating IP address:28443/web** format.

The username and password of the MRS cluster need to be entered for accessing clusters with Kerberos authentication enabled, for example, user **admin**. They are not required for accessing normal clusters with Kerberos authentication disabled.

For the first access, add the site to the trusted site list as prompted to continue to open the page.

**Step 13** When logging out of Manager, terminate and close the SSH tunnel.

----End

## 8.3.6 How Do I Handle Abnormal Status of Core Nodes in an MRS Cluster After Successful Expansion?

### Symptom

Cores nodes are added, but some instances on the nodes may fail to be started. The symptoms are as follows:

1. A core node has been added and is displayed on the **Nodes** page.

<input type="checkbox"/>	Node Name/Resource ID	IP
<input type="checkbox"/>	2222ZSnM0001 1911b9db-94d1-4cbe-8eb1-eeb3088ac146	192.168.1.62

- Some tasks for adding nodes fail or are partially successful.

Add host ❗ Failed 62% Mar 30, 2024

Total Records: 88  < 1 2 3 4 5 6 ... 9 >

### Task List


Name	Status	Pro...
Run decommissioning/recomm...	✔ Successful	100%
Add host	❗ Partially Successful	100%

- If IAM users have been synchronized, you can view unstarted roles on the **Components** page.
- If they are not synchronized, you can view unstarted roles on the Manager page of this cluster.

## Procedure

Scenario 1: The task for adding nodes fails before component installation.

**Step 1** Perform the following steps if the MRS cluster is a pay-per-use cluster:

- Log in to the MRS console.
- Choose **Active Clusters** and click the name of a cluster to access its details page.
- Click  in the upper part of the page. In the **Task List** column, click the task for adding core nodes.
- Records all nodes in the verification request parameter.
- Click the **Nodes** tab, select the nodes recorded in [Step 1.4](#), click **Stop** in the upper right corner, and stop the nodes as prompted.
- Reduce nodes by referring to [Scaling In a Cluster](#).

**Step 2** If the MRS cluster is billed on a yearly/monthly basis, unsubscribe from the abnormal nodes by referring to [Unsubscribing from a Specified Node in a Yearly/Monthly Cluster](#).

----End

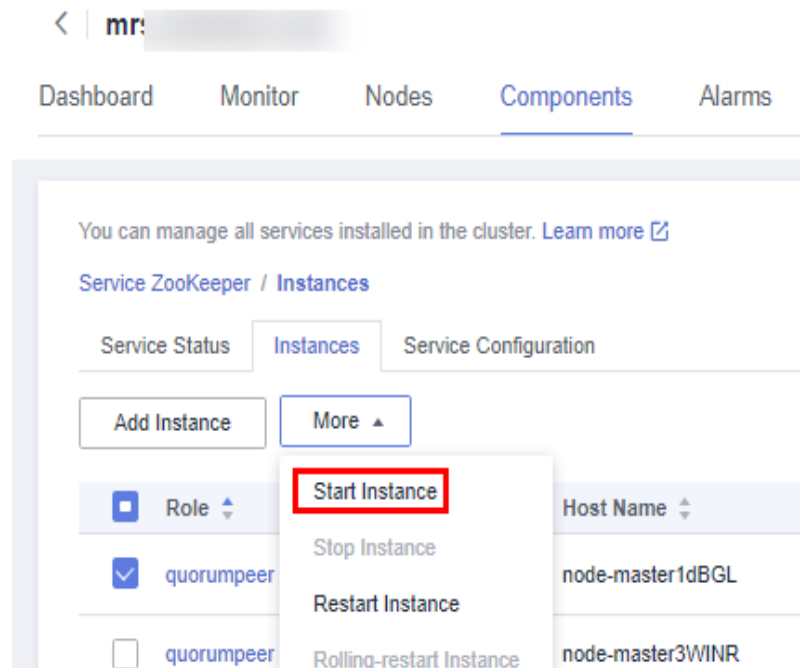
Scenario 2: The task for adding nodes fails after component installation.

**Step 1** Log in to the MRS console.

**Step 2** Choose **Active Clusters** and click the name of a cluster to access its details page.

**Step 3** On the **Dashboard** tab, click **Synchronize** next to **IAM User Sync** to synchronize IAM users.

**Step 4** Click **Components** and check the role status of each service. If a role is not started, select the role, click **More**, and select **Start Instance** to start the instance.



**Step 5** If the startup fails, rectify the fault based on the error information in the task list and try again.

**NOTE**

- If there are many abnormal roles, click **Management Operations** in the upper right corner to start all components.
- For other exceptions that cannot be resolved, contact technical support.
- You can also start the instance on Manager. For details, see [Managing MRS Role Instances](#).

----End